


INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁵ : G11C 7/00		A1	(11) International Publication Number: WO 92/00590
			(43) International Publication Date: 9 January 1992 (09.01.92)
(21) International Application Number: PCT/US91/04484 (22) International Filing Date: 24 June 1991 (24.06.91) (30) Priority data: 546,071 27 June 1990 (27.06.90) US 678,914 1 April 1991 (01.04.91) US 678,912 1 April 1991 (01.04.91) US (71) Applicant: MOS ELECTRONICS CORPORATION [US/ US]; 914 West Maude Avenue, Sunnyvale, CA 94086 (US). (72) Inventor: CHAN, Alfred, K. ; 109 Jacklin Circle, Milpitas, CA 95035 (US). (74) Agent: KIVLIN, B., Noël; Skjervén, Morrill, MacPherson, Franklin & Friel, 25 Metro Drive, Suite 700, San Jose, CA 95110 (US).		(81) Designated States: AT, AT (European patent), AU, BB, BE (European patent), BF (OAPI patent), BG, BJ (OAPI patent), BR, CA, CF (OAPI patent), CG (OAPI patent), CH, CH (European patent), CI (OAPI patent), CM (OAPI patent), CS, DE, DE (European patent), DK, DK (European patent), ES, ES (European patent), FI, FR (European patent), GA (OAPI patent), GB, GB (Eu- ropean patent), GN (OAPI patent), GR (European pa- tent), HU, IT (European patent), JP, KP, KR, LK, LU, LU (European patent), MC, MG, ML (OAPI patent), MN, MR (OAPI patent), MW, NL, NL (European pa- tent), NO, PL, RO, SD, SE, SE (European patent), SN (OAPI patent), SU, TD (OAPI patent), TG (OAPI pa- tent). Published <i>With international search report.</i> <i>Before the expiration of the time limit for amending the</i> <i>claims and to be republished in the event of the receipt of</i> <i>amendments.</i>	

(54) Title RANDOM ACCESS CACHE MEMORY



(57) Abstract

A memory cache apparatus (72) compatible with a wide variety of bus transfer types including non-burst and burst transfers. In burst mode, a "demand word first" wrapped around quad fetch order is supported. The cache memory system decouples the main memory subsystem from the host data bus so as to accommodate parallel cache-hit and system memory transfer operations for increased system speed and to hide system memory write-back cycles from the microprocessor. Differences in the speed of the local and system buses are accommodated, and an easy migration path from non-burst mode microprocessor based systems to burst mode microprocessor based systems is provided. The memory cache apparatus (72) comprises a random access memory (72A-72D), a host port (HP), and a system port (SP). The memory cache apparatus (72) further comprises an input latch connected to the host port for selectively writing data to the memory and an output register connected to the system port for receiving data from the memory and selectively furnishing the data to the host port or the system port. In one embodiment, the input latch is a memory write register, and the output register comprises a read hold register and a write back register.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AT	Austria	ES	Spain	MG	Madagascar
AU	Australia	FI	Finland	ML	Mali
BB	Barbados	FR	France	MN	Mongolia
BE	Belgium	GA	Gabon	MR	Mauritania
BF	Burkina Faso	GB	United Kingdom	MW	Malawi
BG	Bulgaria	GN	Guinea	NL	Netherlands
BJ	Benin	GR	Greece	NO	Norway
BR	Brazil	HU	Hungary	PL	Poland
CA	Canada	IT	Italy	RO	Romania
CF	Central African Republic	JP	Japan	SD	Sudan
CG	Congo	KP	Democratic People's Republic of Korea	SE	Sweden
CH	Switzerland	KR	Republic of Korea	SN	Senegal
CI	Côte d'Ivoire	LI	Liechtenstein	SU	Soviet Union
CM	Cameroon	LK	Sri Lanka	TD	Chad
CS	Czechoslovakia	LU	Luxembourg	TG	Togo
DE	Germany	MC	Monaco	US	United States of America
DK	Denmark				

- 1 -

RANDOM ACCESS CACHE MEMORY

CROSS REFERENCE TO RELATED APPLICATION

This application is a continuation-in-part of co-
5 pending, commonly assigned U.S. patent application Serial
No. 07/546,071 filed June 27, 1990. The above-referenced
application is incorporated herein by reference in its
entirety.

BACKGROUND OF THE INVENTION10 Field of the Invention

This invention relates to memory subsystems, and more
specifically to random access cache memory systems.

DESCRIPTION OF THE RELEVANT ART

Various bus transfer mechanisms are used by present
15 day microprocessors. The bus transfer mechanisms for two
particularly popular microprocessors, the model 80386 and
model 80486 microprocessors available from Intel Corpora-
tion, Santa Clara, California, are summarized below.
Further detail is contained in various publications
20 available from Intel Corporation, including the 386(Tm) DX
Microprocessor Data Sheet, November 1989, and the i486(Tm)
Microprocessor Data Sheet, April 1989.

In the model 80386 microprocessor, a complete data
transfer to or from memory occurs during what is known as
25 a "bus cycle." A bus cycle includes at least two "bus
states;" a bus state is the shortest time unit of bus
activity and requires one processor clock period.
Additional bus states added to a single bus cycle are
known as "wait states." The model 80386 microprocessor
30 may be provided with either 16-bit or 32-bit wide
memories. In a 16-bit system, each group of 16 bits is
considered to be a physical word, and begins at an address
that is a multiple of 2. In a 32-bit wide system, each
group of 32 bits is considered to be a physical

- 2 -

"doubleword," and begins at a byte address that is a multiple of 4. Memory addressing is flexible, and accommodates the transfer of, for example, a logical operand that spans more than one physical doubleword or one physical word, or that is a doubleword operand and begins at an address not evenly divisible by 4, or that is a word operand split between two physical doublewords. Dynamic data bus sizing is supported. The model 80386 microprocessor has separate, parallel buses for data and address. The data bus is 32-bits in width and is bidirectional. The address bus provides a 32-bit value using thirty signals for the thirty upper-order address bits, and four byte-enable signals to indicate the active bytes.

Many of the bus transfer features of the model 80386 microprocessor bus were provided in the model 80486 microprocessor bus. Some of these basic features are illustrated in Figures 1 and 2. Figure 1 shows basic two clock, no wait state, single read and write cycles. The first cycle, a write cycle comprising bus states 1 and 2, is initiated when the address status signal ADS# is asserted at an edge of clock signal CLK in bus state 1. At this time, signal A2-A31 provides a valid address to the system memory; at a later time in bus state 1, signal D0-D31 makes available valid data to the system memory. When the system memory accepts data in accordance with write/read signal W/R#, the external system asserts the ready signal RDY#. In Figure 1, this occurs at the end of bus state 2. The second cycle, a read cycle comprising bus states 3 and 4, is initiated when the address status signal ADS# is asserted at an edge of clock signal CLK in bus state 3. At this time, signal A2-A31 provides a valid address to the system memory. When the system memory returns data in accordance with write/read signal W/R#, the external system asserts the ready signal RDY#. In Figure 1, this occurs at the end of bus state 4.

Figure 2 shows the use of wait states. Wait states are used because the bus cycle time of many commercially

- 3 -

available microprocessors is much shorter than the read/write time required by the conventional low cost DRAM memory generally used as system memory. A faster microprocessor must "wait" for the system memory to complete its read or write, which is accomplished by the insertion of one or more wait states into the bus cycle. For example, the second cycle of Figure 2, which is a write cycle, includes three bus states 5, 6 and 7. Bus state 5 is analogous to bus state 1 of Figure 1, while bus state 7 is analogous to bus state 2 of Figure 1. Bus state 6 is a wait state, inserted because the ready signal READY# was not asserted until bus state 7. Additional wait states are asserted if necessary. The address and bus cycle definition remain valid during all wait states. Similarly, the third bus cycle, a read cycle, includes three bus states 8, 9 and 10, bus state 9 being a wait state.

The model 80486 microprocessor provides a number of additional features, including an internal cache, a burst bus mechanism for high-speed internal cache fills, and four write buffers to enhance the performance of consecutive writes to memory. Accordingly, the model 80486 microprocessor supports not only single and multiple non-burst, non-cacheable cycles, but also single and multiple burst or cacheable cycles.

Burst memory access is used to transfer data rapidly in response to bus requests that require more than a single data cycle. During a burst cycle, a new data item is strobed into the microprocessor every clock. The fastest burst cycle (no wait state) requires two clocks for the first data item (one clock for the address, one clock for the corresponding data item), with subsequent data items returned from sequential addresses on every subsequent clock. Note that in non-burst cycles, data is strobed at best in every other clock.

Burst mode operation is illustrated in Figure 3. A burst cycle, a burst read in Figure 3, begins with an

- 4 -

address being driven and signal ADS# being asserted during the first bus state 12, just as in a non-burst cycle. During the four subsequent bus states, four data items 15, 17, 19 and 21 are returned. Note that during a burst 5 cycle, ADS# is driven only with the first address. The addresses of the data occur within the same 16-byte aligned area, so that external hardware is able to calculate the addresses of the subsequent transfers in advance of the next bus state. For a word size of 32 10 bits, for example, the cache line is four words. The burst mode is indicated when burst ready signal BRDY# is driven active and signal RDY# is driven inactive at the end of each bus state 14, 16, 18 and 20 in the burst cycle. The external memory is signaled to end the burst when the last 15 burst signal BLAST# is driven active at the end of the last bus state 20 in the burst cycle.

Cache memory systems have been developed to permit the efficient use of low cost, high capacity DRAM memory. Cache memory subsystems store recently used information 20 locally in a small, fast memory. When bus transfers are limited to the microprocessor-cache data path, system speed increases.

Cache memory may be internal to the microprocessor, as in the model 80486 microprocessor of the Intel 25 Corporation, or external. An example of an external cache memory for a typical computer system is illustrated in Figure 4. Microprocessor 22 is connected to a local address bus 24 and a local data bus 26. Similarly, cache memory 28 is connected to the local address bus 24 and the 30 local data bus 26. If cache controller 30 determines that data corresponding to the address requested is resident in the cache memory 28, the data is transferred over the local data bus 26. If cache controller 30 determines that the data is not resident in the cache memory 28, the 35 address and data are transferred through cache bus buffer/latch 32A and 32B respectively to the system bus 34. Cache bus buffer/latch 32A and 32B are controlled by

- 5 -

cache bus controller 36, which receives control information from cache controller 30. System memory 38 and system peripherals 40 are connected to the system bus 34.

5 The concept of "cacheable cycles" may be understood with reference to Figure 5, which shows signals typically used in connection with the internal cache of the model 80486 microprocessor of Intel Corporation. Four data items are read from the high speed internal cache memory
10 to the microprocessor in eight clocks without wait states. A cycle is initiated when the address status signal ADS# is asserted during bus state 42. Bus state 42 involves a cache fill, as established by activation of the cache enable signal KEN#. The signal BLAST# remains inactive
15 during bus state 42. The first cycle terminates with the data transfer to the processor in bus state 44. Three additional data cycles consisting of, respectively, bus states 46 and 48, bus states 50 and 52, and bus states 54 and 56, are needed to complete the cache fill. Signal
20 BLAST# remains inactive until the last transfer in the cache line fill, which occurs in bus state 56.

Cache memory has been used for burst transfers, as shown in Figure 3. A cache fill is indicated when the signal KEN# is activate and the signal BLAST# is inactive
25 during bus state 12. The signal BLAST# remains unknown in successive bus states 14, 16 and 18 and is activated only in the fourth successive bus state 20, so that four data items may be burst in succession. The external system informs the microprocessor that it will burst the line in
30 by driving signal BRDY# active during the four successive bus states 14, 16, 18 and 20 in which data is transferred.

At some point after a cache write hit, main memory must be updated. The most widely used methods of updating main memory are write-through and write-back. In write-
35 through, main memory is automatically updated at the same time the cache is written. The processor must wait until the write is completed before it may resume execution. A

- 6 -

variation, usually called "posted" write-through, uses a buffer into which the write data is latched while the processor continues execution. The latched data is then written to main memory whenever the system bus is
5 available. In a write-back cache, new data written to cache is not passed on to main memory until a replacement cycle occurs.

SUMMARY OF THE INVENTION

The cache memory system according to the present
10 invention is compatible with a wide variety of bus transfer types, including non-burst and burst transfers. In burst mode, a "demand word first" wrapped around quad fetch order is supported.

The cache memory system of the present invention
15 decouples the main memory subsystem from the host data bus, so as to accommodate parallel cache-hit and system memory transfer operations for increased system speed, and to hide system memory write-back cycles from the microprocessor. Differences in the speed of the local and
20 system buses are accommodated, and an easy migration path from non-burst mode microprocessor based systems to burst mode microprocessor based systems is provided. In addition, various memory organizations are accommodated, including direct-mapped or one-way set associative, two-
25 way set associative, and four-way set associative.

These and other advantages are achieved in the present invention, in accordance with which a memory cache apparatus comprises a random access memory, a host port, and a system port. The memory cache apparatus further
30 comprises an input latch connected to the host port for selectively writing data to the memory and an output register connected to the system port for receiving data from the memory and selectively furnishing the data to the host port or to the system port. In one embodiment, the
35 input latch is a memory write register, and the output register comprises a read hold register for furnishing

- 7 -

data to the host port and a write back register.

In accordance with another aspect of the invention, a method is provided for operating a memory cache apparatus wherein the memory cache apparatus includes a random access memory, a host port, a system port, an input register coupled to the host port, and an output register coupled to the system port. The method comprises the steps of latching input data into the input register from the host port, comparing a received address to a plurality of cache addresses, loading replaced data from the random access memory into the output register if the received address does not match one of the plurality of cache addresses, loading the input data into the random access memory, and providing the replaced data to the system port.

In accordance with still a further aspect of the present invention, a computer system comprises a host microprocessor having a host address bus and a host data bus, a system memory having a system address bus and a system data bus and a dual port cache memory having a system port connected to the system data bus and a host port connected to the host data bus. A cache controller is further connected to the cache memory.

The invention will be more readily understood by reference to the drawings and the detailed description. As will be appreciated by one skilled in the art, the invention is applicable to cache memory systems in general, and is not limited to the specific embodiment disclosed.

30 BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 (Prior Art) is a set of waveforms illustrating two clock, no wait state, single read and write cycle bus transfer features of the 80386 and 80486 microprocessors.

35 Figure 2 (Prior Art) is a set of waveforms illustrating the use of wait states in basic bus transfer

- 8 -

features of the 80386 and 80486 microprocessors.

Figure 3 (Prior Art) is a set of waveforms illustrating a burst read operation of the 80486 microprocessor.

5 Figure 4 (Prior Art) is a block diagram illustration of an external cache memory for a typical computer system.

Figure 5 (Prior Art) is a set of waveforms typically used in connection with the internal cache of the 80486 microprocessor.

10 Figure 6 is a block diagram of a computer system based on the model 80386 microprocessor.

Figure 7 is a block diagram of a computer system based on the model 80486 microprocessor.

Figure 8 is a block diagram of the data path of a
15 burst RAM cache memory in accordance with the present invention.

Figures 9 and 9A are diagrams illustrating a direct-mapped, one-way associative cache having one bank.

Figures 10 and 10A are diagrams illustrating a two-
20 way set associative cache having two banks.

Figure 11 is a diagram illustrating four burst RAM cache memory chips arranged in a 32-bit configuration.

Figures 12A and 12B are diagrams illustrating generalized data paths and registers within cache
25 memory 72.

Figure 13 is a diagram illustrating the host interface between the 486 microprocessor and cache controller 70 and cache memory 72.

Figure 14 is a diagram illustrating the system
30 interface between a system bus and cache memory 72 and controller 70.

Figure 15 is a diagram showing the internal organization of the memory update registers set 116.

Figure 16 is a set of waveforms showing control and
35 data signals for a single host port read operation.

Figure 17 is a set of waveforms showing control and data signals for a host port burst read operation.

- 9 -

Figure 18 is a set of diagrams showing control and data signals for a host port single write operation.

Figure 19 is a set of waveforms showing control and data signals for a system port single read operation.

5 Figure 20 is a set of waveforms showing control and data signals for a system port single write operation.

Figure 21 is a set of waveforms showing control and data signals for a buffered host to system bypass operation with update.

10 Figure 22 is a set of waveforms showing control and data signals for a buffered host to system bypass operation without update.

Figure 23 is a set of waveforms showing control and data signals for a system to host port bypass operation.

15 Figure 24 is a set of waveforms showing control and data signals for a system to host port bypass operation with reordering.

Figure 25 shows an example cache line used in the operation of the system to host port bypass sequence of
20 Figure 26.

Figure 26 is a set of waveforms showing control and data signals for a system to host port bypass operation with update and partially dirty line.

Figure 27 is a set of waveforms showing control and
25 data signals for an advance write and subsequent quad-fetch operation.

Figure 28 is a set of waveforms showing control and data signals for a read tag miss operation with one write-back.

30 Figure 29 is a set of waveforms showing an advance write operation with subsequent quad fetch and one write-back.

Figure 30 is a set of waveforms showing control and data lines for an advance write operation with subsequent
35 quad fetch and one write-back to a neighboring line.

Figure 31 is a set of waveforms showing control and data signals for several operations occurring within the

- 10 -

burst RAM cache memory.

Figure 32 is a diagram showing internal blocks of cache controller 70.

Figure 33 is a diagram showing functional block pin 5 groups of cache controller 70.

Figure 34 is a set of waveforms showing control and data signals for a system read cycle operation.

Figure 35 is a set of waveforms showing control and data signals for a system write cycle operation.

10 Figure 36 is a set of waveforms showing control and data signals for a buffered NCA write cycle operation.

Figure 37 is a set of waveforms showing control and data signals for a controller register read operation.

15 Figure 38 is a set of waveforms showing control and data signals for a controller register write operation.

Figure 39 is a set of waveforms showing control and data signals for a 486 CPU burst read cache hit operation.

20 Figure 40 is a set of waveforms showing control and data signals for a 486 CPU non-burst read cache hit operation.

Figure 41 is a set of waveforms showing control and data signals for a read line miss and resulting quad fetch operation.

25 Figure 42 is a set of waveforms showing control and data signals for a read line miss operation with no replacement.

Figure 43 is a set of waveforms showing control and data signals for a read line miss operation with reordering.

30 Figure 44 is a set of waveforms showing control and data signals for a multiple write-back cycle operation.

Figure 45 is a set of waveforms showing control and data signals for a cacheable write hit cycle operation.

35 Figure 46 is a set of waveforms showing control and data signals for a write tag miss operation with one write-back cycle and concurrent processing.

Figure 47 is a set of waveforms showing control and

- 11 -

data signals for a write line miss operation and resulting system quad fetch.

Figure 48 is a set of waveforms showing control and data signals for a write tag miss operation with one 5 write-back cycle.

Figure 49 is a set of waveforms showing control signals for the optional address transceivers.

Figure 50 is a set of waveforms showing control and data signals for a flush activation operation, followed by 10 acquiring local bus and first write-back.

Figure 51 is a set of waveforms showing control and data signals for a snoop read miss operation.

Figure 52 is a set of waveforms showing control and data signals for a snoop read hit operation.

Figure 53 is a state diagram illustrating the initial 15 sequencing of the concurrent bus control unit of cache controller 70.

Figure 54 is a state diagram illustrating sequencing during a read tag miss operation.

Figure 55 is a state diagram illustrating sequencing 20 during a write tag miss operation.

Figure 56 is a state diagram illustrating sequencing during a read line miss and write line miss operation.

Figure 57 is a diagram showing the state machines 25 within bus controller 200.

Figure 58 is a diagram showing the state machines within bus controller 202.

DESCRIPTION OF THE EMBODIMENTS

The following includes a detailed description of the 30 best presently contemplated mode for carrying out the invention. The description is intended to be merely illustrative of the invention and should not be taken in a limiting sense.

A burst RAM cache memory in accordance with the 35 present invention is suitable for use with a variety of microprocessors. A block diagram of a computer system

- 12 -

based on the model 80386 microprocessor 60 is shown in Figure 6. The cache memory 72 consists of four substantially identical byte-wide burst RAM memory ICs 72A, 72B, 72C and 72D to support the 32-bit (4-byte) bus system. Address information is transferred between the microprocessor 60, the cache controller 70, and cache memory 72 over local address bus 24, while data is transferred between the microprocessor 60, the cache controller 70, and the host port HP of the cache memory 72 over the local data bus 26. Control signals are communicated between microprocessor 60 and cache controller 70 over control line 64, and control signals are communicated to the cache memory 72 over control line 66. The cache memory 72 is also provided with a system port SP, which is connected over a bidirectional multiple signal line 74 to the system bus 34. The cache controller 70 is provided with a system bus control port CSB, which is connected over a bidirectional multiple signal line 76 to the system bus 34.

A block diagram of a computer system based on the model 80486 microprocessor 62 is shown in Figure 7. The system illustrated in Figure 7 is similar to that of Figure 6 with the exception of the address bus connected to microprocessor 62 which is, in Figure 7, bidirectional. During normal operation, address bus 27 is driven by microprocessor 62. The address bus 27 is driven through the system bus, during a cache invalidation cycle.

Referring next to Figure 8 (comprising Figures 8A, 8B, and 8C), a diagram of the burst RAM cache memory chip in accordance with the invention is shown. The burst RAM memory chip is illustrative of each of the burst RAM memory chips 72A-72D of Figures 6 and 7. The three major sections of the burst RAM memory chip as shown in Figure 8 are the RAM array section 100, the address latches and multiplexer section 102, and the control logic and transceivers section 104. The burst RAM memory chip is also shown with control signal lines within each section

- 13 -

which illustrate some of the control signals for control of the data paths. The details of how the control signals affect data flow are discussed below.

The organization of RAM array section 100 is first considered. The RAM array section 100 is organized as a two-way set associative cache without data buffers, and includes two banks 106 and 108 of 2k x 36 bit static random access memory ("SRAM") (plus parity bits). The two bank division readily accommodates for either a two-way set associative cache or a 64K byte direct-mapped cache. In addition, more burst RAM memory chips can be added for larger caches.

Each bank 106 and 108 of RAM array section 100 is divided into four subarrays I-IV. Each subarray includes 2k x 8 bit memory locations and further includes a parity bit associated with each 8-bit memory location.

An address bus 174 is connected to and provides addressing signals to each of banks 106 and 108. The addressing signals are decoded within each of the banks 106 and 108 to thereby select one of the 2k locations within each subarray I-IV.

As noted, RAM array section 100 is suitable for many memory organizations, including a direct mapped cache organization and a two-way set associative organization. The concepts needed to understand these various memory organizations are illustrated in Figures 9, 9A, 10, and 10A, respectively. Associativity refers to the number of banks of the cache into which a memory block may be mapped. A bank, also known as a frame, is the basic unit into which a cache memory is divided. A direct-mapped (one-way associative) cache such as that shown at 60 in Figure 9 has one bank 61. A two-way set associative cache such as that shown at 70 in Figure 10 has two banks 71A and 71B. A bank is equal to the cache size in a direct-mapped cache and one-half the cache size in a two-way set associative cache.

A "page," which corresponds in size to a cache bank,

- 14 -

is the basic unit into which the physical address space is divided. For example, in Figures 9 and 10, page 0 is shown at 63A in physical memory 62, at 73A in physical memory 72, and at 83A in physical memory 82 respectively.

5 Both memory banks and "pages" are subdivided into blocks; a block is the basic unit of cache addressing. Blocks are shown at 64A and 64B, and at 74A and 74B in Figures 9 and 10 respectively.

Cache address information is stored in a directory.

10 For the direct-mapped cache of Figure 9, the single directory 65 includes 2048 cache address entries, each with three bit fields. The first bit field is a 12-bit directory tag for selecting one of the 2 EXP 12 pages of main memory. The second bit field is an 12-bit set
15 address for selecting one of the 2048 sets in the cache. The third bit field is a 3-bit line address for selecting one of eight lines in a set.

For the two-way associative cache of Figure 10, each of the two directories 75A and 75B includes 1024 cache
20 address entries, each with three bit fields. The first bit field is a 17-bit directory tag for selecting one of the 2 EXP 17 pages of main memory. The second bit field is a 10-bit set address for selecting one of the 1024 sets in the cache. The third bit field is a 3-bit set address
25 for selecting one of eight lines in a set.

The term "directory tag" used above refers to that part of a directory entry containing the memory page address from which that particular block was copied. Any block with the same offset within a page may be mapped to
30 the same offset location in a bank. The tag identifies the page from which the block came.

The term "set" used above refers to all of the directory entries associated with a particular block offset. The number of sets equals the number of blocks. In
35 a two-way set associative organization, a set has two entries, each pointing to a different bank.

The term "line" used above refers to the basic unit

- 15 -

of data transferred between the cache and main memory. A block consists of contiguous lines. Each line in a block has a corresponding "line valid" bit in the block directory entry. Line valid bits are set when the line is written and are cleared when the block tag changes. The number of lines in a block and the line size is determined by the number of valid bits in each directory entry or by cache controller convention. For example, although each directory entry for the cache memory chip of Figure 8 could have enough valid bits to give each 32-bit doubleword in the block a valid bit, the associated cache controller instead operates on four doublewords as a line. A "hit" or "miss" decision is based on the presence or absence of a line within the cache. It is noted that the cache controller 70 associated with the cache memory chip of Figure 8 designates two lines for each tag, and hence, there are eight thirty-two bit doublewords per tag.

Organization of the cache in direct-mapped mode is shown in Figure 9A. Each physical address that the CPU asserts can map into only one location in the cache. The address of each cycle driven by the CPU is broken down into several components: the set index, tag, line select, and doubleword select. CPU address lines A15:5 (11 bits) compose the set index, and determine which cache location the address can map into. Address lines A27:16 (12 bits) are the tag. Assertion of a bus cycle by the CPU generates a tag comparison, with A27:16 being compared against the tag for the given set index. A match of all 12 bits indicates a cache hit. Address line A4 selects between the two lines of a block. Address lines A3 and A2 select individual doublewords within a line, and are not included in cache hit/miss determinations.

Note that Figure 9A is drawn abstractly for understanding of the cache organization. The tag array bits, tag valid, doubleword valid, and the doubleword dirty bits are contained in the cache controller 70, while the actual data for doublewords 0-3 are contained in the

- 16 -

cache memory 72.

Organization of the cache in 2-way set associative mode is shown in Figure 10A. As opposed to the direct-mapped strategy described above, each physical address can map into two locations in the cache. As a result, the 2048 entries are mapped into two parallel sets of 1024 entries each.

Correspondingly, the cache set index is now 10 bits instead of 11, and the tags are 13 bits instead of 12. A bus cycle now triggers two comparison operations, one for each set, at the given set index, with a hit (match) possibly occurring in either set. The results of the comparison are OR'd together, with a high value from the OR output indicating a tag hit.

It will be appreciated that the size of each cache bank, the number of banks, the number of blocks, and the number of lines per tag may be varied without departing from the spirit and scope of the invention.

Referring back to Figure 8, consider next the address latches and multiplexer section 102. Address signal ADDR carried on the host address bus 24 is an 11-bit (<14:4>) signal that addresses the RAM array 100.

An address multiplexer 103 and two address registers, hit address register 109 and miss address register 110, are included in section 102. Under certain circumstances, bits <14:4> of the address signal ADDR are latched into the hit address register 109 and furnished therefrom through the address multiplexer 103 to RAM array 100. The address signal ADDR from the CPU 60 is latched into hit address register 109 through activation of the ADS# signal asserted by the CPU 60. The address information latched in hit address register 109 is provided to the address decoder of the RAM array 100.

Under other circumstances, the 11-bit output of the hit address register 109 is latched into the miss address register 110 and furnished therefrom through the address multiplexer 103 to RAM array 100. Under still other

- 17 -

circumstances, the 11-bit address signal ADDR is furnished directly through the address multiplexer 103 to RAM array 100. These circumstances are described below.

Hit address register 109 includes a control line for receiving a signal CALE or the ADS# signal, and miss address register 110 includes a control line for receiving a signal MALE. A high assertion of the MALE signal causes latching of the address signal into miss address register 110.

Miss address register 110 is used during miss cycles to latch the initial miss address. This address is later used during miss processing, as explained below, so that data retrieved from the system can be correctly updated into the RAM array section 100. It is noted that the address in the hit address register 109 is latched into the miss address register 110 through the activation of signal MALE.

Consider next the control logic and transceivers section 104. The burst RAM cache memory chip is provided with a 9-bit (including parity) host port HP 113 which is connected to the local (host) data bus (bus 26 of Figures 6 and 7) through a 9-bit line. Each burst RAM chip is also provided with a 9-bit (including parity) system port SP 112 which transfers an appropriate byte of information between the system data bus (bus 34 of Figures 6 and 7) through its 9-signal line. It should be noted that the local and system data buses are four bytes or 32-bits wide, so thus four burst RAM chips are required to support the 32-bit systems as shown in Figures 6 and 7. For one implementation, as illustrated in Figure 11, burst RAM cache memory chip 72A supports the most significant byte of the local and system buses, burst RAM cache memory chip 72B the next most significant byte, burst RAM cache memory chip 72C the next most significant byte, and burst RAM cache memory chip 72D the least significant byte. A single doubleword (where each word comprises sixteen data bits) is stored in four bytes. Each cache memory chip

- 18 -

72A-72D stores one of the four bytes of a doubleword. Each doubleword is stored within a single set of the subarrays labelled either I, II, III, or IV. A "line" refers to the four adjacent doublewords stored within each set of the subarrays I-IV. The set of four subarrays labeled I of burst RAM chips 72A-72D contain the 32-bit doubleword located at the line's highest address. Subarray sets II, III, and IV contain the third, second, and first addressed doublewords in the line, respectively.

Referring back to Figure 8, the burst RAM chip data path further includes three sets of holding registers 114A-114D, 116A-116D and 118A-118D, and a memory write register 120. A memory read hold register set ("MRHREG") includes four 8-bit registers 114A-114D and is provided to support data bus burst read operations on the host data port 113. Each of the four registers 114A-114D includes an additional bit for parity. A memory write back register set ("MWBREG") includes four 8-bit (plus parity) registers 118A-118D and is provided to accommodate burst write operations on the system data port 112. A memory update register set ("memory update register set 116") includes four 8-bit (plus parity) registers 116A-116D and is provided to accommodate quad fetch miss data operations from system memory 38. Finally, memory write register 120 is an 8-bit register (plus parity) provided to accommodate scalar write operations on the host data port 113.

Figures 12A and 12B are shown to simplify the conceptual architecture of the burst RAM cache memory 72 in accordance with the invention. The architecture shown in Figures 12A and 12B corresponds to that of Figure 8, and includes hit address register 109, miss address register 110, RAM array section 100, read hold register set 114, memory update register set 116, write back register set 118, and write register 120. The generalized diagrams of Figures 12A and 12B may be referred to for simplifying the descriptions of the cache memory 72 contained herein.

- 19 -

The burst RAM cache memory 72 and controller 70, in accordance with the invention, control the paths of data in response to various control signals. These signals, along with the terminal pins of cache memory 72 and controller 70 from which and to which they are provided, are listed below in Tables I and II with a brief description of their purposes. In addition, Figures 13 and 14 illustrate the system and host interface connection pins for an 80486 computer system incorporating a cache memory 72 and controller 70 in accordance with the invention.

TABLE I**CONTROLLER 70 SIGNALS****LOCAL PROCESSOR INTERFACE**

<u>Signal</u>	<u>Type</u>	<u>Description</u>
CCLK	Input	External Clock Input. This pin is directly connected to the i486 CPU CLK pin. CCLK provides the fundamental timing and internal operating frequency for the controller 70. CCLK only needs TTL levels for proper operation. All external timing parameters are referenced to the rising edge of CCLK.
A<31:2>	Input/ Output	Processor Local Physical Address. As inputs, these lines provide the physical memory and I/O addresses of the local bus for the controller 70. As outputs, addresses are driven back to the 486 processor during cache invalidation cycles. These signals directly connect to 486 local processor address bus.

- 20 -

CONTROLLER 70 SIGNALS

BE<3:0>#	Inputs	<p>Byte Enables. These four lines determine which bytes must be driven valid for read and write cycles to external memory.</p> <p>BE3# associates with D<31:24> BE2# associates with D<23:16> BE1# associates with D<15:8> BE0# associates with D<7:0></p> <p>These four signals directly connect to 486 BE<3:0>#</p>
Reset	Input	<p>RESET Input. The RESET input is asynchronous. The activation of the RESET pin clears all tag valid bits in the controller 70 tag directory array together with all the other internal states.</p>
A20M#	Input	<p>Address 20 Mask. Asserting the A20M# input causes the controller 70 to mask physical address bit A<20> before performing a tag directory comparison and before driving a memory cycle to the outside world.</p> <p>When A20M# is asserted active, the controller 70 emulates the 1 Mbyte address space of the 8086. The signal is only asserted when the host CPU is in real mode.</p> <p>This signal is connected to the A20GATE signal of most IBM PC/AT compatible chipsets.</p>

PROCESSOR CYCLE DEFINITIONS

M/IO#	Input	<p>Processor Memory/IO Access. Distinguishes between memory or I/O cycles. This signal directly connects to the 486 M/IO# pin.</p>
D/C#	Input	<p>Processor Data/Code Access. Distinguishes between data or code cycles. This signal is directly connected to the 486 D/C# pin.</p>
W/R#	Input	<p>Process Read/Write Access. Distinguishes between read or write cycles. This signal is directly connected to the 486 W/R# pin.</p>

- 21 -

CONTROLLER 70 SIGNALS

LOCK#	Input	Bus Lock. Assertion of this signal indicates the 486 needs to have exclusive right to the local bus for a read-modify-write operation. This signal directly connects to the LOCK# pin of 486.
--------------	--------------	---

READY INDICATION AND LOCAL BUS STATES

ADS#	Input	Address Status Output. This signal indicates that the address and bus cycle definition signals are valid. ADS# is active on the first clock of a bus cycle and goes inactive in the second and subsequent clocks of the bus cycle. This input has a weak pull-up resistor.
-------------	--------------	---

The controller 70 uses ADS# together with other ready indication signals to monitor 486 local bus activity.

- 22 -

CONTROLLER 70 SIGNALS

RDY#	Input	<p>Local Bus Cycle Ready Input. In 486 systems, assertion of RDY# indicates the completion of any local (Weitek) bus cycles. This signal is ignored on the end of the first clock of a bus cycle. This input has a weak pull-up resistor.</p> <p>A low assertion of RDY# will cause the RDYO# output from the controller 70 to go low.</p> <p>A programmable option through the controller 70 allows asynchronous RDY# input. This asynchronous option allows a coprocessor with slow output delay to interface with the controller 70. In asynchronous mode, the controller 70 will forward RDY# to the RDYO# output in the next clock. In synchronous mode, RDY# will be forwarded to RDYO# in the same clock provided that setup time is met. After reset, RDY# is assumed to be synchronous.</p> <p>This signal is directly connected to the 486 RDY# pin.</p>
------	-------	--

- 23 -

CONTROLLER 70 SIGNALS

RDYO#	Output	<p>System Non-Burstable Cycle Ready Output. The controller 70 completes cycles such as cache read hits, write cycles (hit/miss) and cycles specifically directed at the controller 70 (control register update cycles). For cycles not directly completed by the controller 70, such as read misses and system cycles, the controller 70 will forward either the SBRDY1# or SRDY1# signals from the system memory bus as RDYO# to the CPU.</p> <p>For cycles to NCA (Non-Cacheable Address) regions, I/O, Halt/Shutdown, INTA cycles, and cycles when the controller 70 is disabled, either SRDY1# or SBRDY1# can be returned by the system. However, either of these signals will be passed to the CPU as RDYO#.</p> <p>This signal is directly connected to the RDY# input pin of the 486. In case a numerics co-processor (NPX) is present in the system READY# of the NPX will be connected to RDY# of the controller 70. The controller 70 will then forward READY# from the NPX to the CPU.</p>
BRDYO#	Output	<p>486 Burst Cycle Ready Output. In 486 systems, the controller 70 drives this output on to the processor local bus indicating the completion of a controller 70 burst read hit data cycle. In cache subsystems using cache memory 72 Burst-RAMs, the controller 70 will forward the SRBRY1# signal from the system memory bus for cache read miss cycles.</p> <p>This signal is directly connected to the 486 BRDY# pin.</p>

- 24 -

CONTROLLER 70 SIGNALS

BLAST#	Input	<p>Burst Last. This signal indicates that the next BRDY# or RDY# returned will terminate the 486 host cycle. The controller 70 only samples BLAST# in the second or subsequent clocks of any bus cycle.</p> <p>This signal is directly connected to the 486 BLAST# pin.</p>
BOFF#	Output	<p>Host CPU Back-Off. This signal is used by the controller 70 to obtain the 486 local bus. During snoop read hits, the controller 70 asserts BOFF# to the 486 one clock after AHOLD is asserted in order to access the cache data array. See section entitled "Snoop Operations" for a complete discussion.</p>
HOLD	Output	<p>Host CPU Bus Hold Request. This signal is used for flush operations in order to obtain the local CPU bus. During either hardware or software flushes, the controller 70 will assert HOLD to the CPU. HOLD is released upon completion of the flush operation.</p>
HLDA	Input	<p>Host CPU Bus Hold Acknowledge. The assertion of HLDA indicates that the controller 70 has been granted the local bus to begin flush operations.</p> <p>Upon recognition of HLDA, the controller 70 will begin generating write-back cycles to the system to clear lines which contain dirty data.</p>
LBA#	Input	<p>Local Bus Access. This pin indicates to the controller 70 that the current bus cycle should occur only on the host (local CPU) bus. Assertion of this signal will prevent any system read or write operations from occurring as a result of the current cycle. However, this signal must be asserted to the controller 70 in the T1 state for proper operation.</p>

CACHE CONTROL

- 25 -

CONTROLLER 70 SIGNALS

FLUSH#	Input	<p>FLUSH#. The FLUSH# input is asynchronous and need only be active for one clock. This input has a weak pull-up resistor.</p> <p>This input clears all valid, dirty and LRU bits in the controller 70. The controller 70 will copy all dirty valid data back to system memory before executing the flush operation.</p> <p>This pin should be directly connected to the FLUSH# pin of the 486.</p>
PCD	Input	<p>Page Cache Disable. This pin provides a cacheable/non-cacheable indication on a page-to-page basis from the 486. The 486 will not perform a cache fill for any data cycle when this signal is asserted.</p> <p>PCD reads are cached by the controller 70. PCD read cycles which are cache hits are treated as normal cacheable cycles, with data being returned to the CPU in zero wait states. However, this data will not be cached inside the CPU. PCD write cycles generate buffered write-through cycles to the system.</p> <p>This signal is connected directly to the 486 PCD pin.</p>
PWT	Input	<p>Page Write Through. Assertion of this signal during a write cycle will cause the controller 70 to treat the current write cycle as a write-through cycle.</p> <p>A hit on a PWT write cycle will cause an update both in the data cache and main memory. A miss on a PWT write cycle will update only main memory, and will not generate a system quad fetch.</p> <p>This signal is directly connected to the 486 PWT pin.</p>

- 26 -

CONTROLLER 70 SIGNALS

KEN#	Output	<p>Cache Enable. KEN# is used to indicate to the 486 if the data returned by the current cycle is cacheable.</p> <p>A cycle to a protected address region will cause the controller 70 to deassert KEN# in T1 to prevent the i486 CPU from performing any cache line fills. KEN# will continue to be deasserted until RDY# or BRDY# is returned, whether the cycle is a cache hit or miss.</p> <p>For all Weitek 4167 cycles, KEN# will be deasserted. Data passing between the 486 and the Weitek co-processor is not cached. Also, in order to support long instruction execution of the Weitek 4167, the controller 70 will continue to deassert KEN# upon detection of a Weitek cycle till RDY# assertion is received.</p> <p>This signal is directly connected to the 486 KEN# input.</p>
-------------	---------------	--

CACHE INVALIDATION CONTROL

EADS#	Output	<p>External Address Valid. EADS# indicates that a valid external address has been driven onto the 486 address pins A<27:4>. These address bits are then checked with the cache tag directory inside the 486. If a match is detected, the directory entry associated with A<27:4> will immediately be invalidated. A<31:28> will always be driven to zero along with A<27:4>.</p> <p>The controller 70 will assert EADS# for main memory write cycles as indicated by MEMWR# signal being high from system memory bus interface.</p> <p>This signal is connected directly with the 486 EADS# pin.</p>
--------------	---------------	--

- 27 -

CONTROLLER 70 SIGNALS

AHOLD	Output	<p>Address Hold Request. Asserting this signal will force the 486 to float the address bus in the next clock. While AHOLD is active, only the address bus will be floated, the data portion of the bus may still be active.</p> <p>The controller 70 uses AHOLD to attain address bus mastership for performing an internal cache invalidation to the 486 when system writes occur.</p> <p>This signal is directly connected to the AHOLD pin of the 486.</p>
-------	--------	---

CACHE MEMORY CONTROL INTERFACE

HPOEA#, HPOEB# (G0#,G1#)	Outputs	<p>Host Port Output Enables. These signals are connected to the host port data output enable inputs of Cache RAMs to individually enable the selected cache bank to drive the data bus. If cache memory 72 is used, these signals will be connected to the G0# and G1# inputs.</p>
HPWEA#, HPWEB# (HW0,HW1)	Output	<p>Host Port Write Enables. These signals are connected to the write enable inputs of the Cache RAMs in order to individually enable the selected cache bank to receive data. If cache memory 72 is used, these signals are connected to the HW0# and HW1# inputs.</p>
CCS<3:0>#	Outputs	<p>Cache Chip Select. These signals are connected to the chip select inputs of the Cache RAMs associated with each byte of the data word. If cache memory 72 is used, these signals are connected to the SELECT# inputs.</p> <p>During all processor write cycles, these outputs emulate BE<3:0># to select those bytes that are updated in main memory during a partial write. In read hits and read miss cache update cycles, all four signals are active.</p>

- 28 -

CONTROLLER 70 SIGNALS

MALE	Output	<p>Miss Address Latch Enable. This signal activates after ADS# activation if either a read/write miss has been detected. Assertion of MALE and HPWEA# or HPWEB# in the same clock will inhibit writing to data RAM array section inside Burst-RAM cache memory 72..</p> <p>Activation of MALE will reset all valid bits associated with each data byte in memory update register set 116.</p>
WBSTB (BMUXC<0>)	Output	<p>Write-Back Strobe. The rising edge of WBSTB results in the Burst-RAM data entry associated with miss address register 110 to be latched into write back register set 118 (write-back register). This data will be written to main memory later if it is dirty. For multiple replacement cycles, the controller 70 will assert WBSTB on the clock after MWB is asserted to allow Burst-RAM data entry associated with miss address register 110 to be latched into write back register set 118.</p>
RHSTB (BMUXC<1>)	Output	<p>Read Hit Strobe. The rising edge of RHSTB results in the Burst-RAM data entry associated with hit address register 109 to be latched into read hold register set 114 (read hold register). This data will be burst to the 486.</p>
QWR	Output	<p>Quad Write. Activation of this signal results in the data in memory update register set 116 (update register) being written into cache memory 72 data array entry pointed to by the address in miss address register 110 and the cache memory 72 internal select logic.</p> <p>This signal is connected to QWR input of cache memory 72.</p>

- 29 -

CONTROLLER 70 SIGNALS

DW# (QWRWQ)	Output	<p>Dirty Word. During a system quad fetch, the controller 70 uses this signal to indicate dirty/non-dirty status for each fetched doubleword to the cache memory 72's. Active (low) assertion of this signal, qualified with SRDYI# or SBRDYI#, indicates that dirty data is currently present in the cache data array, and should not be overwritten by the fetched data.</p> <p>This signal is directly connected to DW# inputs of the cache memory 72's.</p> <p>DW# is also used as a signal to the system memory controller during quad write cycles, to indicate dirty/non-dirty status of the corresponding doubleword. In the non-dirty case, the system can ignore the driven data and immediately return SRDYI# or SBRDYI#.</p>
BYPASS	Output	<p>Host/System Bypass. This signal connects host port data to system port data of the cache memory 72. Cache miss reads and non-cacheable read/writes will activate the BYPASS signal.</p> <p>This signal is connected directly to the BYPASS input of the cache memory 72.</p>
SP OE# (SW#)	Output	<p>System Port Output Enable. A low assertion of this signal enables the Burst-RAM system port outputs. For cache memory 72 writes, SP OE# will be asserted for all four data transactions.</p> <p>During write cycles, this signal is asserted one clock after SADS# to be compatible with 486 write cycles.</p> <p>This signal is directly connected to the SP OE# pin of the cache memory 72.</p>

- 30 -

CONTROLLER 70 SIGNALS

MWB	Output	<p>Multiple Write-Back. MWB is used during write-back cycles, in the case when both lines associated with a replaced tag are dirty. MWB is asserted by the controller 70 at the end of the first write-back cycle, in order to toggle miss address register 110 to point at the second line corresponding to the replaced tag entry.</p> <p>This signal is connected directly to the MWB input of the cache memory 72.</p>
HA<3:2>	Output	<p>Host Port Address <3:2>. These two bits indicate the word address within a quad word. These bits are part of the address bit associated with data at the host port.</p> <p>These two signals connect directly to HA<3:2> inputs of the cache memory 72.</p>
CALE (HALE)	Output	<p>Controller ALE. CALE is generated by the controller 70 during the first bus state of controller-initiated cycles. The controller 70 will generate CALE during flush operations and snoop read hits. Upon assertion of CALE, the cache memory 72 will latch the controller 70-generated address.</p> <p>This pin is connected directly to the cache memory 72 CALE pins.</p>

SYSTEM BUS INTERFACE

SADS#	Output	<p>System Bus Address Status. SADS# is the system equivalent of 486 ADS#. This indicates that valid address SA<27:2>, SBE<3:0># and cycle definition signals (SM/IO#, SW/R#, SD/C#) are available to the system.</p> <p>SADS# is asserted in the first clock of a system bus cycle. This pin is tri-stated when the controller 70 does not have system bus ownership.</p>
-------	--------	---

- 31 -

CONTROLLER 70 SIGNALS

SM/IO#	Output	System Bus Memory/IO. Distinguishes system bus memory or I/O cycle accesses. This pin is tri-stated when the controller 70 does not have system bus ownership.
SD/C#	Output	System Bus Data/Code Access. Distinguishes system bus data or code accesses. This pin is tri-stated when the controller 70 does not have system bus ownership.
SW/R#	Output	System Bus Write/Read Access. Distinguishes system bus write or read accesses. This pin is tri-stated when the controller 70 does not have system bus ownership.
SA<27:4>	I/O	System Address Bus. The controller 70 uses these inputs as snoop system bus address when some other system bus master controls the bus.

The controller 70 will drive these address lines to system memory during miss processing, system and write-through cycles.

The controller 70 will float SA<27:4> during system idle states or at the end of system bus cycles. SA<27:21> have weak pull-down resistors.

SA<3:2>	I/O	System Bus Address. These signals indicate the address of each 32-bit doubleword within a quad line. They play a similar role as SA<27:4> except during burst cycles, SA<3:2> are wrapped around a line (16 byte) boundary.
---------	-----	---

The controller 70 can control SA<3:2> sequentially or use the 486 burst-order. This is controlled by a bit in the control register.

Cache memory 72 Burst-RAMs will connect these two signals to SA<3:2>.

To facilitate easy interface to system memory, SA<3:2> timing is earlier than SA<27:4>.

- 32 -

CONTROLLER 70 SIGNALS

SBE<3:0>#	I/O	System Bus Byte Enables. As inputs, these pins are used during snoop write operations to select the proper bytes to be written into the cache data array. As outputs, these pins indicate to the system which bytes are active for system read/write operations.
SRDYI#	Input	System Bus Non-Burst Cycle Ready Input. Assertion of SRDYI# indicates the completion of any non-burst system bus cycles. Simultaneously asserting both SRDYI# and SBRDYI# signals will pass SRDYI# to the CPU. The controller 70 will forward only SRDYI# input to RDYO# for non-cacheable cycles. This input has a weak pull-up resistor.
SBRDYI#	Input	System Bus Burst Cycle Ready Input. Assertion of this signal indicates that the current system bus burst cycle is completed. The controller 70 will ignore SBRDYI# assertion at the end of the first clock of a system bus cycle. This input has a weak pull-up resistor. The controller 70 will pass SBRDYI# back to the host CPU as BRDYO#, except for system read cycles. For these cycles, assertion of either SRDYI# or SBRDYI# will be passed to the CPU as RDYO#. SBRDYI# for Halt/Shutdown, I/O, and INTA cycles will also be passed to the CPU as RDYO#.

- 33 -

CONTROLLER 70 SIGNALS

SBLAST#	Output	<p>System Burst Last. This signal asserted (low) indicates that the next assertion of SBRDYI# or SRDYI# will terminate the system bus cycle. SBLAST# is asserted for all system bus cycles.</p> <p>As an enhancement over the 486, SBLAST# is driven to a valid level in T1, instead of being indeterminate. Hence, SBLAST# will be valid in the same clock as SADS#.</p> <p>SBLAST# will be asserted during the last transfer of any system bus cycle.</p> <p>This pin is tri-stated during system bus hold.</p>
SLOCK#	Output	<p>System Lock. This signal is passed from the local CPU bus to the system bus. LOCK# is asserted by the CPU for indivisible read-modify-write operations. Assertion of LOCK# will trigger assertion of SLOCK#, indicating to system logic that the 486/controller 70 should retain the bus until LOCK# is deasserted.</p>

SYSTEM DATA BUS TRANSCEIVER CONTROL

ST/R#	Output	<p>System Data Bus Data Transmit/Receive. This signal defines the direction of the optional system bus data transceivers.</p> <p>This signal is connected to the DIR pin of the 646 data transceivers.</p>
SD OE#	Output	<p>System Data Output Enable. This enables the output of the optional data transceivers. The controller 70 will deassert SD OE# if another slave device on the system bus is granted bus ownership.</p> <p>This signal is connected to OE# of external 646 data transceivers.</p>

- 34 -

CONTROLLER 70 SIGNALS

SLDSTB Output Local Data Strobe. The rising edge of this signal latches data into the optional external 646 data transceivers.

In quad fetch mode, the controller 70 will assert SLDSTB four times.

SYSTEM ADDRESS BUS TRANSCEIVER CONTROL

SA CP Output System Address Clock Pulse. The rising edge of this signal latches data into the latch transceiver that drives SA<27:2>, SBE<3:0>#, SM/IO#, SW/R# and SD/C# in the system bus.

The rising edge of this signal latches the address and cycle definition signals from the controller 70.

SACP will be asserted only at the beginning of a burst cycle.

This signal is connected to the external latch transceiver CAB pin.

SA OE# Output System Address Output Enable. This signal enables the external address latch or external address latch transceiver outputs when the controller 70 is the current bus master and disables them otherwise.

The controller 70 will deassert SA OE# if SHOLD is acknowledged.

This signal is connected to OE# of the external address bus latch or external address bus latch/transceivers.

SA DIR Output System Address Direction. This signal controls the DIR (direction) input of the optional address transceivers. SA DIR is high when the controller 70 owns the bus and is low when the controller 70 grants ownership of the system bus. SA DIR toggles low or high the clock following change of SHLDA.

SYSTEM BUS ARBITRATION SIGNALS

- 35 -

CONTROLLER 70 SIGNALS

SHOLD	Input	System Bus Hold Request. This signal is used to request system bus mastership from another slave device. The role of this signal is equivalent to HOLD in the processor local bus.
SHLDA	Output	System Bus Hold Acknowledge. This signal is used to indicate that a request for system bus ownership has been granted. During SHLDA assertion, SDOE# and SAOE# will be deasserted.

SYSTEM BUS CACHE INVALIDATION REQUEST

SEADS#	Input	System External Address. SEADS# indicates that a valid system bus address has been driven onto the system memory bus SA<27:2>. A match of this address will invalidate the cache directory entry inside the controller 70. This input has a weak pull-up resistor. Unlike the i486 CPU, the controller 70 provides for partially valid lines. To support this, SA<3:2> must be driven by the system to correct levels along with SA<27:4> when SEADS# is asserted.
--------	-------	---

- 36 -

CONTROLLER 70 SIGNALS

Snoop writes: SEADS# and SMEMW/R# high should be asserted by the system during main memory writes. The controller 70 will forward SEADS# and the invalidate address to the CPU. A <31:28> will be driven as zeroes on the processor local bus to properly invalidate the internal CPU cache. The controller 70 will assert EADS# for only one clock for each system bus snoop write. In addition, the SBE<3:0># signals and write data should be driven to correct levels. This data will be updated into the cache data array if a snoop write hit occurs.

Snoop reads: Assertion of SEADS# and a low assertion of SMEMW/R# will result in the controller 70 asserting SMEMDIS if an address match has been detected in the controller 70 Cache directory. The controller 70 will write dirty data associated with the driver address onto the system bus for correct operation. No EADS# assertion will be sent to the local processor bus for a system bus read cycle.

The controller 70 will sample SEADS# every clock.

COHERENCY SUPPORT (see also SHOLD, SEADS#) SA<27:2>

SMEMW/R#	Input	System Bus Memory Write/Read Cycle. This signal indicates whether a snoop read or snoop write is occurring.
		Snoop reads trigger SNPBUSY and a tag lookup, while snoop writes trigger an invalidate to the 486 CPU.
SMEMDIS	Output	System Bus Memory Access Disable. This signal is asserted when the controller 70 detects a snoop read hit so that dirty data can be sent from the cache memory 72 Burst-RAMS.

- 37 -

CONTROLLER 70 SIGNALS

SNPBUSY	Output	<p>Snoopbusy. This signal goes high in response to a snoop cycle.</p> <p>For snoop writes, the falling edge of SNPBUSY indicates that the memory controller should return SRDYI# or SBRDYI# to complete the snoop write operation.</p> <p>For snoop reads, SNPBUSY remaining high indicates that the necessary tag lookup has not yet completed. On the falling edge of SNPBUSY, the state of SMEMDIS indicates whether the controller 70 or the memory system will supply data to satisfy the snoop read request.</p>
----------------	--------	--

CONTROL REGISTER INTERFACE

MCCSEL#	Input	<p>controller 70 Select. Assertion of this signal indicates that the current local bus cycle is addressed to the controller 70. These cycles include controller 70 Control Register Address Index load cycles and Control Register data read and write cycles.</p> <p>MCCSEL# is connected to local bus decode logic output for the controller 70 address cycles. This input has a weak pull-up resistor.</p>
D<8:0>	Input/ Output	<p>Cache Control Data Bus. These nine data pins are connected to the least significant byte of local processor data bus. They are used to load the controller 70 Control Register index address and data content from the 486.</p> <p>In diagnostic mode, D<8:0> can be used as an output to read the internal states of the controller 70.</p>

MULTIPROCESSOR SUPPORT (see also SEADS#, SHOLD and SHLDA)

- 38 -

CONTROLLER 70 SIGNALS

SBREQ	Output	Pending System Bus Request. The controller 70 asserts SBREQ if an internal system request is pending. SBREQ is asserted the same time as SADS#. In case the controller 70 does not currently own the bus, SBREQ will be asserted the same clock that SADS# would have, had the controller 70 owned the bus.
-------	--------	---

SBREQ will never be floated.

This signal plays the same role as 486 BREQ pin in local processor bus.

TABLE II**CACHE MEMORY 72 SIGNALS**

<u>Signal</u>	<u>Type</u>	<u>Description</u>
BURST RAM INPUT LINES		
ADDR<14:4>	Inputs	486 Local Address Bus. These bits address one of 2048 entries of 32-bit double-words in each burst-RAM data array.
ADS#	Input	Address/Data Strobe. This is the latch enable strobe for CPU-generated bus cycles. The falling edge of this signal creates a flow-through mode for the hit address register 109. The rising edge of ADS# latches the address into hit address register 109.
CALE (HALE)	Input	Controller Address Latch Enable. This is the latch enable strobe for controller-generated bus cycles. CALE opens the hit address register 109 level latch to ADDR.

- 39 -

CACHE MEMORY 72 SIGNALS

MALE	Input	<p>Miss Address Latch Enable. This signal should be asserted after ADS# or CALE activation if a miss for either read or write cycles has been detected. MALE performs several functions. First, it latches the address in hit address register 109 into miss address register 110. It also inhibits any host port writes from being performed on clock edges that it is sampled high.</p> <p>In addition, MALE latches bank select information in order to direct a subsequent update into the cache array. Activation of MALE with either HPOEA# or HPWEA# indicates that the next quad write (QWR) operation will update bank A, while MALE and either HPOEB# or HPWEB# selects bank B. Bank selection, as described above, must be qualified with assertion of the SELECT# input.</p> <p>MALE has a higher priority than does the MWB input. Assertion of both MALE and MWB on the same clock edge will result in recognition of MALE, but not MWB.</p>
RESET	Input	<p>Cache memory 72 Reset. Reset should be active for at least four CLK clocks for cache memory 72 to complete reset. Reset will clear all mask and valid bits.</p>
WBSTB (BMUXC<0>)	Input	<p>Write Back Strobe. The rising edge of WBSTB results in the cache memory 72 data entry associated with miss address register 110 being latched into write back register set 118. This data should be written to main memory if any is dirty.</p>
RHSTB (BMUXC<1>)	Input	<p>Read Hit Strobe. The rising edge of RHSTB results in cache memory 72 data entry associated with hit address register 109 to be latched into read hold register set 114. This data will subsequently be burst to the 486 CPU.</p>
HP<8:0>	Inputs/ Outputs	<p>Byte-wide (with parity) data input/output to and from the host CPU.</p>

- 40 -

CACHE MEMORY 72 SIGNALS

DW# (QWRWQ)	Input	Dirty Word. Active assertion of this signal qualifies the miss fetch word associated with either SRDYI# or SBRDYI# to be written into the data array of the cache memory 72.
HA<3:2>	Inputs	These two bits select a doubleword within a given line. These bits are part of the address associated with the data at the host port of the cache memory 72's.
HPWEA#, HPWEB#, (HWO,HW1)	Inputs	Host Port Write Enables. HPWEA# is the host port write enable for bank A and HPWEB# is the host port write enable for bank B. A low assertion of either HPWEA# or HPWEB# indicates the corresponding bus cycle generated by the host is a write cycle. Either of these two signals being sampled low on a clock edge will trigger the latching of HP data into write register 120. These two signals cannot be both active simultaneously.
HPOEA#, HPOEB#, (G0#,G1#)	Inputs	Host Port Buffer Output Enables. HPOEA# will enable data from bank A for the corresponding cycle, while HPOEB# will enable data from bank B. Both signals cannot be active simultaneously. Along with HPOEA# and HPOEB#, SELECT# must be asserted in order to enable the host port outputs.
SP<8:0>	Inputs/ Outputs	Byte-wide (with parity) data input/output to and from main memory.
SA<3:2>	Inputs	These two bits select a doubleword within a given line. These address bits are part of the address associated with the data at the system port of the cache memory 72.
SRDYI#	Input	A low activation of this signal indicates that the data read from main memory is ready. The low level of this signal, qualified with a clock edge, triggers the data at the system port data pins to be latched into memory update register set 116 in the burst-RAM at the doubleword location by SA<3:2>.

- 41 -

CACHE MEMORY 72 SIGNALS

SBRDYI#	Input	Assertion (low) of this signal indicates that data is available on the system port. The low level of this signal, qualified with the clock, latches the system port data pins into memory update register set 116 at the doubleword location selected by SA<3:2>.
SPOE# (SW#)	Input	<p>System Port Output Enable. A low assertion of SPOE# indicates that the system port will be performing a write operation to main memory. For burst writes, SPOE# should be asserted for all four write data transactions. Along with SPOE#, SELECT# must be asserted to enable the system port outputs.</p> <p>For Bypass operations, SPOE# acts as a direction control signal. BYPASS asserted high with SPOE# low creates a host-to-system port bypass, with the contents of write register 120 being driven onto the system port. BYPASS asserted while SPOE# is high will generate a system-to-host bypass, with system port data being passed directly to the host port.</p>
QWR	Input	Quad Write. Activation of Quad Write results in the data residing in memory update register set 116 being written into the cache memory 72 data array entry, at the address pointed to by miss address register 110 and the cache memory 72 internal bank select logic. QWR overrides the mux control logic. In addition, QWR resets both the bank select information previously latched through assertion of MALE, and all mask and valid bits associated with memory update register set 116.
BYPASS	Input	This signal connects the host port data HP<8:0> to system port data pins. Specifically, assertion of BYPASS and SPOE# high connect the system port data pins to the host port. To achieve a host-to-system bypass, BYPASS asserted high and SPOE# asserted low connect the output of write register 120 to the system port data pins.

- 42 -

CACHE MEMORY 72 SIGNALS

SELECT#	Input	Cache memory 72 Chip Select. This signal asserted (low) indicates that the corresponding cache memory 72 is selected. A high assertion of this signal will result in all data output pins being tri-stated. For any data outputs to be enabled, SELECT# should be asserted. For complex operations which utilize both the host and system ports, SELECT# should be asserted for the entire operation.
MWB	Input	Multiple Write-Back. MWB should be used for write-back cache architectures where each tag entry corresponds to two lines. MWB is asserted during write-back cycles, in the case where the other line associated with the replaced tag is dirty and needs to be written back to memory. The assertion (high) of MWB toggles the A4 bit of the address stored in miss address register 110. Subsequent assertion of WBSTB then loads write back register set 118 with the second line of data to be written back to the system.

- 43 -

Referring back to Figures 8 and 12A, all writes to the RAM array section 100 from the host port 113 propagate through the memory write register 120. In addition, the memory write register 120 can be used as a buffer on write 5 cycles, for either buffered write-through cycles or buffered non-cacheable writes. For example, write operations to addresses defined as non-cacheable addresses are buffered in memory write register 120 and consequently can be completed in zero wait states on the local data bus 10 26. These writes continue on the system data bus 34 until the system accepts the written data. At the same time, local bus operations may continue. This operation is explained in further detail below.

The read hold registers 114A-114D are used to allow 15 one clock-burst read operation from the cache memory 72. During the first transfer of a burst read, 32 bits of data are read into read hold registers 114A-114D. To complete the second, third and fourth transfers, the contents of the read hold registers 114A-114D are driven on the local 20 CPU data bus 26, one doubleword at a time. A burst read hit causes all four 32-bit doublewords within the same line to be read into the read hold registers 114A-114D.

"Wrapped-around" burst order is further supported. The first demand doubleword is fetched and sent to the 25 host port 113 directly from the RAM array section 100. The subsequent three remaining words are next fetched through the read hold registers 114A-114D. This architecture allows very high speed (50 MHz and beyond) burst mode operation without using ultra-fast (sub 10 ns) 30 data RAMs. Burst-order is controlled by signal HA<3:2> and is totally transparent to the cache memory 72.

The write-back registers 118A-118D are used to hide the write-back cycles that occur when system data from read misses replace dirty data. The older line, 35 containing the dirty data, is latched into the write-back registers 118A-118D, allowing the system read to occur without delay. At the completion of the system fetch, the

- 44 -

contents of the write-back registers 118A-118D are written back to the system bus 134. For best performance, the cache memory 72 and controller 70 in accordance with the invention allow these quad writes to be burst to the
5 system, if the system memory can accept such bursts.

For miss operations in write-back mode, the write-back registers 118A-118D hold the data from the selected data line in the RAM array section 100 to be replaced. This data is written back to main memory if the replaced
10 line contains valid and dirty data. This allows burst writes on the system memory port 112 without requiring access to the RAM array section 110 for write-back replacement cycles. Therefore, the RAM array section 100 is available to serve the host port 113 for local bus read
15 and write hits.

Memory update registers 116A-116D are used as a holding register for incoming data from system quad fetches due to read and write miss cycles. As each doubleword is returned by the system, it is passed on to
20 the CPU 60 and latched inside one of the memory update registers 116A-116D. At the completion of the system fetch, all four doublewords are written into the RAM array section 100 by assertion of the quad write (QWR) signal.

Inclusion of the memory update registers 116A-116D in
25 the burst RAM architecture allows systems to gain a performance benefit. Since incoming system port 112 data is latched into the memory update registers 116A-116D instead of the RAM array section 100, the array is free to service local bus hit operations as the read miss
30 processing occurs on the system port 112.

As explained above, the memory update registers 116A-116D contain quad fetch miss data from main memory. The order of loading the memory update registers 116A-116D is controlled through signal SA<3:2>, making burst-order to
35 main memory purely transparent to the burst RAM. Once the line is loaded into the memory update registers 116A-116D, the entire line is loaded into the cache RAM array section

- 45 -

100 in one clock by activating the QWR signal.

Associated with each byte of the memory update registers 116A-116D is both a valid and a mask bit. The functionality of the mask and valid bits is shown below.

5 When signal QWR is asserted (high) on a clock edge, the contents of the memory update registers 116A-116D are updated into the RAM array section 100, as pointed to by the address data stored in miss address register 110 and the previously latched bank select information (explained
10 below). Each byte of memory update registers 116A-116D is written to the RAM array section 100 if its valid bit is set (indicating valid data from the system port 112), and its MASK bit is cleared (indicating no advance write occurred for that byte).

15	Condition	On QWR and CLK:
	MASK# & VALID#	NO WRITE
	MASK# & VALID	WRITE OCCURS
	MASK & VALID#	NO WRITE
	MASK & VALID	NO WRITE

20 Figure 15 shows the internal organization of the memory update registers 116A-116D.

Mask bits are set during advance writes by assertion of MALE#, SELECT#, and either HPWEA# or HPWEB#. The mask bit set within memory update registers 116A-116D is
25 selected by SA<3:2>. Setting of the mask bits will be further described below. The valid bits are set by assertion of SELECT# and either SRDYI# or SBRDYI#. As with the mask bits, the valid bit set within memory update registers 116A-116D is selected by SA<3:2>. A further
30 discussion of the valid bits is also given below.

Assertion of QWR (quad writes) clears all valid and mask bits associated with memory update registers 116A-116D. This allows subsequent advance write or system fetch to begin.

35 The bank section of RAM array section 100 is next considered. Host port 113 reads from and writes to bank

- 46 -

106 are selected by HPOEA# and HPWEA#, respectively, and host port 113 reads from and writes to bank 108 are selected by HPOEB# and HPWEB#, respectively. A low assertion of signal HPWEA# signifies a write cycle from the host port 113 to bank 106 through write register 120, and a low assertion of HPWEB# signifies a write cycle from the host port 113 to bank 108 through write register 120. HPOEA# and HPOEB# are active low output enables. A low assertion of HPOEA# will gate the read data from bank 106 to host port 113 and the low assertion of HPOEB# will gate the read data from bank 108 to the host port 113. HPWEA#, HPWEB#, HPOEA# and HPOEB# cannot be active simultaneously. For system port 112 reads and writes, the burst RAM cache memory 72 uses previously-latched inputs for bank selection.

Detection of a CPU miss, indicated by assertion of signal MALE (Miss Address Latch Enable), causes latching of bank select information. After signal MALE is asserted, the bank select information is latched within the burst RAM. Subsequent system port read (RAM array section 100 to write back) and write (memory update register set 116A-116D to RAM array section 100) operations are directed to the bank previously selected when MALE was asserted. QWR (quad write) activation or assertion of signal RESET clears the latched bank select information. Similarly, the SELECT# signal is latched through MALE for the system port miss processing operations between the burst RAM and main memory.

Referring again to Figure 8, the control logic and transceivers section 104 further comprises four multiplexers 111A-111D for routing data, a plurality of control signal decoders 119, 122, 124, 128, 129, 150, 158 and 166, and a plurality of data path drivers designated with solid triangular symbols. Each of the data path drivers is singly identified and further described below.

The operation of the burst RAM chip is next explained for several general and specific operating modes and

- 47 -

cycles as explained in the sections below. Figures 6, 7, and 8 will be referenced in the following description of each mode. Furthermore, within the following description, signal timing diagrams are referenced to illustrate control of the various data paths. Several of the signals shown control decoders 119, 122, 124, 128, 129, 150, and 166. Table III below lists each of these decoders and the functions implemented at their respective output lines.

TABLE III

10	DECODER OUTPUT	FUNCTION IMPLEMENTED
	119	= $\overline{SELECT\#} \cdot BG \cdot (\overline{BYPASS} \cdot \overline{SPOE\#})$
	BG	= $HPOEA\# + HPOEB\#$
	122A	= $SA3 \cdot SA2 \cdot \overline{BYPASS} \cdot \overline{SPOE\#} \cdot SPSEL$
	122B	= $SA3 \cdot SA2 \cdot \overline{BYPASS} \cdot \overline{SPOE\#} \cdot SPSEL$
15	122C	= $SA3 \cdot SA2 \cdot \overline{BYPASS} \cdot \overline{SPOE\#} \cdot SPSEL$
	122D	= $SA3 \cdot SA2 \cdot \overline{BYPASS} \cdot \overline{SPOE\#} \cdot SPSEL$
	122E	= $\overline{BYPASS} \cdot \overline{SPOE\#} \cdot SPSEL$
	SPSEL	= $SPSELP\# + (\overline{SELECT\#} \cdot \overline{BYPASS}) + \overline{SPOE\#}$
	SPSELP	= $(SPSELP\# + \overline{SPSELP\#})$
20	SPSELP	= $\overline{SELECT\#} \cdot \overline{MALE} + MPINC \cdot \overline{MALE}$
	SPSELP	= $(\overline{BYPASS} \cdot \overline{SW\#} \cdot \overline{SELECT\#})$
		+ $\overline{BRESET} + \overline{SELECT\#}$
	128A	= $HA <3> \cdot HA <2>$
	128B	= $\overline{HA} <3> \cdot \overline{HA} <2>$
25	128C	= $HA <3> \cdot \overline{HA} <2>$
	128D	= $\overline{HA} <3> \cdot HA <2>$
	128E	= $RHSTB \cdot \overline{QWR}$
	128F	= $\overline{WBSTB} + \overline{QWR}$
	150A	= $HA3 \cdot HA2 \cdot \overline{QWR} \cdot BHW \cdot (\overline{BHW} \cdot \overline{MALE})$
30	150B	= $HA3 \cdot \overline{HA2} \cdot \overline{QWR} \cdot BHW \cdot (\overline{BHW} \cdot \overline{MALE})$
	150C	= $\overline{HA3} \cdot HA2 \cdot \overline{QWR} \cdot BHW \cdot (\overline{BHW} \cdot \overline{MALE})$
	150D	= $\overline{HA3} \cdot \overline{HA2} \cdot \overline{QWR} \cdot BHW \cdot (\overline{BHW} \cdot \overline{MALE})$
	BHW	= $HPWEA\# + HPWEB\#$
	163	= $\overline{BYPASS} \cdot \overline{SPOE\#} \cdot \overline{SELECT\#}$
35	166	= \overline{QWR}

- 48 -

158D	=	BHW • BWED • BK0SEL
159D	=	BHW • BWED • BK1SEL
BWED	=	BWEBQWR + BWEDHWR
BWEDQWR	=	QWR • VALID0 • MASK0
5 BWEDHWR	=	HA3 • HA2 • BHW • QWR • BYPASS • MALE
BHW	=	HWO# + HWI#
BK1SEL	=	SELECT BANK 1 (108)
BK0SEL	=	SELECT BANK0 (106)
10 158C	=	BHW • BWEC • BK0SEL
159C	=	BHW • BWEC • BK1SEL
BWEC	=	BWECQWR + BWECHWR
BWECQWR	=	QWR • VALID1 • MASK1
BWECHWR	=	HA3 • HA2 • BHW • QWR • BYPASS • MALE
15 158B	=	BHW • BWEB • BK0SEL
159B	=	BHW • BWEB • BK1SEL
BWEB	=	BWEBQWR + BWEBHWR
BWEBQWR	=	QWR • VALID2 • MASK2
BWEBHWR	=	HA3 • HA2 • BHW • QWR • BYPASS • MALE
20		
158A	=	BHW • BWEA • BK0SEL
159A	=	BHW • BWEA • BK1SEL
BWEA	=	BWEAQWR + BWEAHWR
BWEAQWR	=	QWR • VALID3 • MASK3
25 BWEAHWR	=	HA3 • HA2 • BHW • QWR • BYPASS • MALE
124	=	[QWR • BHW + BHW • MALE] • BK0SEL
127	=	[QWR • BHW + BHW • MALE] • BK1SEL

Host Port Operations

30 Host Port Single Reads

Cache memory 72 allows single read operations with the local bus processor through the host port. A single read operation is shown in Figure 16.

- 49 -

The falling edge of ADS# activates the hit address register 109 into flow-through mode. Activation of SELECT# enables the internal bank select logic of the cache memory 72. The falling edge of HPOEA# or HPOEB# indicates to the cache memory 72 which bank will supply the data. The RAM array section 100 is accessed and read to the local processor through the host port 113. After access times of valid address to valid data delay, data valid delay from HPOEx#, and data valid from HA<3:2> have passed, valid data is available on the host port 113 of the cache memory 72.

Host Port Burst Reads

The architecture of the cache memory 72 supports burst mode read operations. Each cache memory chip 72A-72D contains the internal 32-bit read hold register set 114A-114D to facilitate high-speed burst read hit operations. Up to four-transfer burst are possible with the architecture of the cache memory 72. Figure 17 shows a four-transfer burst operation.

The first transfer of a burst read is accessed similarly to the scalar read previously described. ADS# assertion will activate hit address register 109 into a 'flow-through' mode. RHSTB should be deasserted (low) to allow data from the cache RAM array section 100 to bypass the read hold register set 114 and be sent to the host port data pins.

After the completion of the first transfer, signal RHSTB can be asserted. This has two effects. First, the entry in the RAM array section 100 pointed to by the hit address register 109 and bank select inputs will be latched into read hold register set 114 as the rising edge of RHSTB occurs. Second, RHSTB being held high (while BYPASS and WBSTB are low) will connect the output of the read hold register set 114 to the host port 113 data pins.

During the access that the read hold register set 114 is loaded, two timing parameters must be met in order to

- 50 -

supply valid data on the host port. Both tx (RHSTB to valid HP data) and ty (HA<3:2> to valid host port data) must be met. After both of these parameters have been satisfied, the read hold register set 114 drive valid data 5 onto the host port 113.

Once the read hold register set 114 has been loaded, the remaining transfers of the burst read come from the read hold register set 114, making very fast burst accesses possible. Burst order are controlled by the 10 HA<3:2> inputs, with the burst order transparent to the MS443's. As HA<3:2> toggles the next burst-order address, valid data from the read hold register set 114 will be available after ty has passed. RHSTB should be held high and WBSTB low to keep the read hold register set 114 15 connected to the host port.

Note that once read hold register set 114 has been loaded, the RAM array section 100 is available for system port operations. In addition, the host port is available for use while miss processing is occurring on the system 20 side.

Host Port Single Writes

The cache memory 72 supports single host port write operations. The falling edge of ADS# sets hit address register 109 into "flow-through" mode. Similarly, 25 assertion of either the two write enables (HPWEA# or HPWEB#) will cause host port data to begin flowing through the write register 120, as well as select which of the two banks of the RAM array section 100 is to be updated. HPWEA# and HPWEB# can trigger the write operation.

30 Sampling either of these inputs asserted (low) on a rising clock edge, with MALE deasserted, will result in the data in write register 120 to be written into the RAM array section 100. HA<3:2> controls which byte within a doubleword in the RAM array section 100 is to be updated.

35 Figure 18 shows a host port write operation to the cache memory 72.

- 51 -

Activation of signal MALE will inhibit write register 120 data from being written into the data array. The write enable signals are sampled on every clock edge. Hence, a write into the RAM array section 100 will occur 5 only on the rising edge of a clock, if MALE is inactive and either HPWEA# or HPWEB# asserted. As will be shown later, buffered writes can be performed by asserting HPWEA# or HPWEB# to latch data into write register 120, and then inhibiting the array write by asserting signal 10 MALE.

Finally, as the write to the RAM array section 100 occurs, the corresponding bit in the memory update register set 116 will be set. This allows 'advance' writes to occur. Advance write operation is explained in 15 more detail below.

System Port Operations

System Port Single Reads

There is no direct path between the RAM array section 100 and the system port 112. Cycles which require the 20 burst-RAMs 72A-72D to supply data on the system port 112 (such as snoop reads) may be accomplished as follows. First, the assertion of either ADS# or CALE allows the address to flow through hit address register 109. MALE should be asserted to latch the request into miss address 25 register 110. The assertion of WBSTB and HPOEA# or HPOEB# latch data into write back register set 118. Finally, with SPOE# asserted, while BYPASS is deasserted, will enable the contents of write back register set 118 onto the system port 112. SA<3:2> will select among the four 30 bytes of write back register set 118. For read cycles, the states of SRDYI# and SBRDYI# are not used. Figure 19 details a single read operation on the system port 112.

System Port Burst Reads

Burst reads from the RAM array section 100 to the

- 52 -

system port 112 may be accomplished. One case where burst reads from the RAM array section 100 to the system port 112 are necessary are for flushing the cache. These reads would begin as detailed above in the single read case.

5 Once write back register set 118 has been loaded with four bytes of data, these bytes may be burst onto the system bus byte-by-byte, with SA<3:2> toggling to select among the four bytes. SPOE# should remain asserted, and BYPASS deasserted, in order to enable the contents of write back
10 register set 118 onto the system port 112.

System Port Single Writes

As before in the system port read case, there is no path to allow direct writing from the system port 112 to the RAM array section 100. Writes from the system port
15 112 must first propagate through the memory update register set 116 before they can be stored in the RAM array section 100. Snoop writes are an example of system port writes cycles.

Before memory update register set 116 can be written
20 into the RAM array section 100, the bank 106 or 108 of the cache memory 72 which is to be written to must be selected. This is done as described above, by assertion (low) of SELECT# and either the HPOEA# or HPOEB# inputs. The assertion of signal MALE latches both the bank select
25 information and the address into miss address register 110.

Memory update register set 116 is loaded by the assertion of either SRDYI# or SBRDYI# on a rising clock edge. The byte of data appearing on the system port 112
30 is latched into memory update register set 116, as selected by the SA<3:2> inputs.

As the 9 bits of data from the system port 112 are updated into memory update register set 116, the corresponding valid bit will be set if signal DW# is
35 inactive (high). This valid bit being set will allow the corresponding byte of memory update register set 116 to be

- 53 -

written into the cache memory RAM array section 100, when signal QWR is later asserted. The presence of valid bits allow abortion of quad fetch on read miss for a subsequent pending write request in write-through mode to access the
5 main memory interface. Valid bits make partial write of a fetched line possible.

Activation of the DW# signal indicates that the cache memory RAM array section 100 contains dirty data at the same address. Correspondingly, activation (low) of signal
10 DW# inhibits dirty and valid miss data from being overwritten by quad fetch miss data by clearing the corresponding valid bit.

Once miss address register 110 is loaded and the bank selected, the data within memory update register set 116
15 may be written into the RAM array section 100. This write is accomplished by assertion of the QWR signal on a clock edge. The bytes of memory update register set 116 which did not receive any writes from the system port 112 will not be updated into the RAM array section 100, as the
20 valid bit for these bytes is cleared. The quad write operation clears all valid and mask bits associated with memory update register set 116.

Figure 20 shows a single write operation to the RAM array 100 through the system port 112.

25 System Port Burst Writes

System port burst write operations execute similar to the above described operation for system port single writes. Since memory update register set 116 is a 32-bit register, up to four bytes may be loaded into memory
30 update register set 116 before its contents are written into the RAM array section 100. Assertion of SRDYI# or SBRDYI# on the clock edge loads the system port 112 data into the memory update register set 116 byte pointed to by SA<3:2>. The four bytes of memory update register set 116
35 can be loaded in as quickly as four clocks.

- 54 -

Dual-Port Operations

Bypass Operations

For optimum performance, the architecture of the cache memory 72 allows bypass operations. Bypass can occur in either direction as described below.

Host to System Port Bypass

Some host bus cycles may be designated as write-through cycles. The architecture of the cache memory 72 supports these cycles.

10 Activation of BYPASS together with signal SPOE# (low) will create a host-to-system bypass condition, with data from the write register 120 being passed directly to the system port 112. The assertion of BYPASS and SPOE# will override and reset the selection information previously
15 latched by the rising edge of MALE. The controller 70 controlling the cache memory 72 must ensure that all miss processing operations are completed at the system port 112 before any bypass write cycles begin.

Updates to the cache RAM array section 100 may occur
20 during bypass operations as previously described in the section on Host Port Writes. The combination of either HPWEA# or HPWEB# asserted, SELECT# asserted, and MALE negated on a rising clock edge will generate a write into the cache RAM array section 100 during a host-to-system
25 port bypass.

Buffered Host to System Port Bypass

Write operations to the cache memory 72 may occur as buffered writes. As described above, the falling edge of either HPWEA# or HPWEB# allows host port 113 data to begin
30 flowing through the write register 120. Once write register 120 has been loaded with valid data, the buffered write may then be accomplished by asserting BYPASS high and SPOE# low. The contents of write register 120 will be driven onto the system port data pins while these two

- 55 -

inputs remain asserted.

Since buffering occurs in the write register 120, buffered write operations may occur whether or not an update into the cache RAM array 100 occurs.

5 Figure 21 shows a cache update due to a host port 113 write, with the write being buffered and continuing on the system port 112, until the system accepts the write data. Buffered write misses will be detailed in the write miss section; however, they proceed identically except for the
10 MALE input. Unlike the BYPASS read case, in BYPASS writes, the state of the MALE input is recognized.

Figure 22 details a buffered write operation, where no cache update occurs. MALE is asserted to inhibit the write operation.

15 System to Host Port Bypass

System-to-host port bypasses may also be generated. During cache read miss cycles, the requested data may be bypassed asynchronously from the host port 113 to the system port 112 in order to minimize the miss penalty and
20 optimize performance. Use of the BYPASS path allows read miss processing to occur as quickly as possible, with no clock latencies between arrival of incoming data at the system port 112 and forwarding of the same data on to the host port 113. Designers should allow for the BYPASS
25 propagation delay from the system port 112 to the host port 113, in addition to the normal CPU read data setup time.

As the data arrives at the system port 112, it may be bypassed directly to the host port 113 by assertion of the
30 BYPASS and SPOE# signals high. Note the dual functionality of SPOE#. When BYPASS is deasserted, SPOE# is used to enable system port data from write back register set 118 onto the system port 112. When BYPASS is asserted, however, SPOE# acts as a direction control for
35 BYPASS.

SRDYI# or SBRDYI# sampled asserted (low) on a rising

- 56 -

clock edge will latch system port data into memory update register set 116. Since memory update register set 116 is a 36-bit register, SA<3:2> will select one of four bytes in memory update register set 116.

5 In addition, cycles which the cache controller 70 designates as non-cacheable may be easily handled with the BYPASS signal. By asserting BYPASS, the requested data will be supplied by the system port 112 rather than the cache RAM array section 100. Figure 23 shows a non-
10 cacheable cycle, with the assertion of BYPASS and SPOE# held high generating a system-to-host bypass. Note that when BYPASS is asserted, MALE becomes a "don't care" input.

As the system-to-host bypass cycle occurs, system
15 port data may be latched into memory update register set 116. This can be accomplished on rising clock edges by assertion (low) of either the SRDYI# or SBRDYI# inputs. The cache memory 72 does not differentiate between these inputs, and they are internally ANDed together. Assertion
20 of either of these on a clock edge will latch system port data into one byte of memory update register set 116, as selected by SA<3:2>, provided data has been valid one setup time previous to the rising edge of the clock.

Assertion of the QWR signal (on a clock edge) will
25 then write the contents of memory update register set 116 into the cache RAM array 100. Any mask bits set in memory update register set 116 will inhibit writes to the corresponding byte in the RAM array section 100.

System to Host Port Bypass with Reordering

30 Memory update register set 116 can act as a buffer register if the burst-order between the 486 host and main memory is different. Memory subsystems using DRAM nibble mode are likely to use sequential burst order, unlike the i486 CPU. As each of the subsequent three doublewords of
35 the burst are read from main memory, they are bypassed to the host port 113 if the burst orders are the same. A

- 57 -

differential in burst order will result in data coming from memory update register set 116. Signal QWR can be asserted once all miss data from memory is updated into the memory update register set 116.

- 5 Figure 24 details a read miss, where reordering occurs between the system and host ports.

System to Host Port Bypass with Partially Dirty Lines

Some cache architectures may contain lines of data which are partially dirty. Figure 25 shows a line which
10 is partially valid and partially dirty. For i486 CPU reads from such a line, the cache memory 72 must supply the data which is dirty, while the system must supply the portion of the line which is not present in the cache. The architecture of the cache memory 72 supports such
15 situations.

System-to-host port bypass cycles may be interrupted by negation of the Bypass signal. When Bypass is negated, the cache memory 72 will supply data as selected by either HPOEA# or HPOEB#, the host port address, and HA<3:2>.
20 Figure 26 shows a CPU read miss, from the line detailed in Figure 25. During the third transfer, the negation of Bypass causes the cache memory 72 to supply the dirty data from its RAM array 100.

Advance Writes

25 The cache memory 72 architecture supports "advance" host port writes in write-back cache architectures. "Advance" write miss processing means that write miss data can be directly updated into the RAM array section 100 in cache memory 72. System fetches at the same address, in
30 order to fill the remainder of the cache line, can occur subsequently and be written into the RAM array section 100 without overwriting the previously stored data.

There is a mask bit for each byte of memory update register set 116. Mask bits can be used to support
35 advance writes. As a write miss (SELECT# and either

- 58 -

HPWEA# or HPWEB#) from the host port 113 (through write register 120) occurs, the mask for the corresponding byte (as selected by HA<3:2>) in memory update register set 116 is set. Any subsequent fetches from the system port 112 will not overwrite the data now written into the RAM array section 100.

For example, consider a host port write to the RAM array section 100 of the least significant byte (HA<3:2>=0). On the rising edge of the clock that the write is triggered in the RAM array section 100 (HPWEA# or HPWEB# is asserted), the mask bit of memory update register set 116 also corresponding to the least significant byte (SA<3:2>=0) is set. A system fetch can then occur to retrieve the remainder of the line (improving the cache hit rate), filling memory update register set 116 as the fetch occurs. When the contents of memory update register set 116 are written to the RAM array 100 by assertion of signal QWR, any memory update register set 116 bytes with set masks will be protected from overwrite.

At reset, all four mask bits associated with each byte of the memory update register set 116 are cleared.

Figure 27 shows an advance write occurring.

Evacuation

CPU miss cycles which bring new data into the cache from the system port 112 will usually replace valid data in the cache. For write-back cache architectures, the replaced data may be dirty. In these cases, the dirty data must be evacuated from the cache RAM array 100 and written to system memory, so as not to be simply overwritten by the incoming data from the system port 112.

The architecture of the cache memory 72 allows for easy evacuation of dirty data into the write back register set 118. In addition, additional performance is obtainable, since read miss processing can occur at full speed, with write back register set 118 being written to

- 59 -

the system after the system fetch completes.

Host Port Read Miss with Evacuation

Figure 28 details a CPU read miss cycle, necessitating a system-to-host port bypass operation with evacuation. ADS# signals the beginning of the cycle. Signal MALE latches bank select information and the address into miss address register 110. Later assertion of WBSTB latches the data which is to be replaced into write back register set 118. Incoming data from the system port 112 is forwarded directly on to the host port 113 through assertion by signal Bypass. SPOE# is held high to correctly enable the Bypass direction, and directly connect the system port 112 to the host port 113. At the end of the system port cycle, QWR writes the data in memory update register set 116 into the RAM array section 100. Finally, the contents of write back register set 118 must be written to the system port 112.

Host Port Advance Write with Evacuation

The rising edge of WBSTB will trigger the latching of data (selected by miss address register 110 and previously latched bank select information) into write back register set 118. A write-back burst sequence should occur if the data replaced is dirty.

The falling edge of ADS# signals the beginning of the cycle. HPWEA# and HPWEB# are asserted in order to select the bank to be written. However, before the write can occur, the dirty data must first be evacuated from the RAM array 100. As such, MALE should be asserted along with HPWEA# or HPWEB# to inhibit the write operation from occurring. The rising edge of WBSTB will trigger the latching of the data from the selected replace line into write back register set 118. A write-back cycle should occur if the data replaced is dirty.

As in the read miss case, memory update register set 116 will be used to hold quad fetch data from main memory.

- 60 -

Miss data will be fetched in a "wrapped-around" fashion with the demand word fetched first. Each byte of the memory update register set 116 is associated with a valid bit. As each byte is updated into memory update register set 116 through SRDYI# assertion, the corresponding valid bit will be set if signal DW# is active. This valid bit will qualify the corresponding byte of the memory update register set 116 to be written into the cache memory 72 data array as QWR is asserted. After each QWR (quad write) cycle, all valid and mask bits associated with memory update register set 116 will be reset.

Figure 29 details a write miss, with evacuation of dirty replaced data and the ensuing system quad fetch. At the end of the quad fetch, the contents of write back register set 118 are written to memory.

Figure 30 shows the same write miss cycle as before; however, at the end of the quad fetch, signal MWB is asserted in order to toggle miss address register 110 to point to the second line belonging to the replaced tag entry. After MWB is asserted, WBSTB loads write back register set 118 with the contents of this line. Finally a write-back cycle to the system of this newly loaded data occurs.

Concurrent Processing

The dual port architecture of the cache memory 72 is one of its most powerful features. The cache memory 72 is capable of processing on the system and host ports concurrently.

There are several instances where concurrent processing is possible. First, buffered write-through cycles can occur through the write register 120. As the write continues on the system bus, the host port 113 can process read and write hits. In addition, system read and write requests can occur in parallel with host port operations. And, as has been previously shown, write-backs of dirty data can occur from write back register set

- 61 -

118 and be hidden from the CPU. While these write-backs occur, local CPU cycles can be satisfied on the host port 113.

Figure 21 shows many of the features of the cache memory 72 in use simultaneously. A CPU write miss occurs on the host port 113. This miss will be written into the RAM array 100 through an 'advance' write. The dirty data in the RAM array section 100 which is to be evacuated is loaded into write back register set 118. To increase the hit rate of the cache, a system quad fetch occurs to fill the remainder of the line that the CPU write updated. This system fetch is transparent to the CPU, and will be stored in the RAM array section 100 without overwriting the 'advance' write. While this fetch completes, the CPU generates another write cycle and a burst read cycle, which are both satisfied on the host port 113.

Summary of Cache Memory 72

Condition	System Port Data Comes From:
20 SPOE# and BYPASS# and SELECT# SPOE# and BYPASS and SELECT#	WBREG WREG
Condition	Host Port Data Comes From:
SELECT# and SPOE and BYPASS	System Port
25 When SELECT# and BYPASS are low, and either HPOEA# or HPOEB# low and:	Host Port Data Comes From:
RHSTB# and WBSTB# RHSTB and WBSTB# WBSTB	Data Array RHREG MUPREG
30 ADS# or CALE: Falling edge makes HADDREG flow-through. Rising edge latches address into HADDREG.	
MALE: Inhibits write operations to cache data array. 35 HADDREG address latched into MADDREG.	

- 62 -

Latches bank select information for future SP operations.

QWR:

- 5 Contents of MUPREG written into data array at
MADDREG.
MUPREG mask and valid bits cleared.

WBSTB:

Latches data at address in MADDREG into WBREG.

HPWEA# or HPWEB#

- 10 The falling edge of either of these makes WREG flow-
through. Either low on a clock (MALE low) latches HP
into WREG. Activates bank select information.

The architecture and operating modes of cache
controller 70 are next considered. The following sections
15 further include the sequencing states of controller 70.

Referring to Figure 32, a system diagram is shown of
the cache controller 70 and cache memory 72 connected to a
CPU 60 and to memory 61. Cache controller 70 is shown
with a bus controller 200, a bus controller 202, an
20 integrated tag array 204, a concurrent bus control unit
206 and data path control unit 208. Bus controller 200
interfaces with CPU 60 through corresponding address and
control lines, and bus controller 202 interfaces with
memory 61 through corresponding address and control lines.
25 Furthermore, data path control unit 208 generates control
signals that are received by cache memory 72.

Referring to Figure 33, controller unit 70 is shown
with specific input and output terminals adapted for a 486
microprocessor-based system. Controller 70 is shown with
30 local processor interface unit 220, processor cache
invalidate control 222, control register interface unit
224, burst RAM interface unit 226, system interface
control unit 228, system address bus control 230, system
data bus control 232, cache coherency control 234 and
35 system bus arbitration unit 236. Each of units 220
through 226 provide and receive signals from the other
system devices.

- 63 -

Register Set and Programming Model

The default configuration of the controller 70 is in write-back mode, configured for operation in a PC-AT environment. Any programming that is necessary will
 5 generally be carried out by the BIOS or operating system as part of the initialization process. Once the controller 70 is initialized, few instances will arise where additional programming is necessary.

The controller 70 contains the following registers.
 10 A control register (CREG) determines the operating modes of the controller. An expansion register (XREG) offers cascade expansion mode to the controller 70 architecture. Eight address registers allow the controller 70 to offer four protected address regions. A protection register
 15 (PREG) defines the operating modes of these four regions. Availability of these regions eliminates the need for high-speed address decode PALs in a system. These registers will be visible to application programmers, for system customization as desired.

20 Registers in the controller 70 are accessed by a two-step register indexing method. The index address of the register to be read or written is written to the low I/O location assigned to the controller 70 (i.e., MCCSEL# asserted and A2 is low). The contents of that register
 25 can be read or written by performing the corresponding read or write I/O cycle to the high I/O location (i.e., MCCSEL# asserted and A2 is high).

The index value corresponding to each controller 70 programmable register is listed below:

30	Index (Hex)	Register	Bits
	00	Control Register	CREG<7:0>
	01	Expansion Register	XREG<7:0>
	10	Protection Register	PREG<7:0>
	20,	1st protected region	PR1S<19:12>
35	21,	start address	PR1S<27:20>
	23,	1st protected region	PR1E<19:12>
	24,	ending address	PR1E<27:20>

- 64 -

	26,	2nd protected region	PRS2<19:12>
	27,	start address	PRS2<27:20>
	29,	2nd protected region	PR2E<19:12>
	2A,	ending address	PR2E<27:20>
5	2C,	3rd protected region	PR3S<19:12>
	2D,	start address	PR3S<27:20>
	2E,	3rd protected region	PR3E<19:12>
	2F	ending address	PR3E<27:20>
10	30,	4th protected region	PR4S<19:12>
	31	start address	PR4S<27:20>
	32,	4th protected region	PR4E<19:12>
	33	ending address	PR4E<27:20>

Programmable Register Definitions

Control Register (Index Address = 00 Hex) - This register identifies the major operating modes of the controller 70. The bit definition of this register is listed as follows:

Bit Position	Name	Definition
CREG<0>	Reserved	This bit should always be set to a value of 1. A zero in this field is not supported.
20 CREG<1>	ARDY	Setting this bit will make the controller 70 RDY# input asynchronous. In asynchronous mode, RDY# input will be forwarded to RDY0# in the next clock. This allows the coprocessor with slow RDY# delay to interface to the controller 70. In synchronous mode, RDY# input will be forwarded to RDY0# in the same clock.
CREG<2>	Reserved	This bit should always be set to a value of 1. A zero in this field is not supported.
CREG<3>	Reserved	This bit should always be set to a value of 1. A zero in this field is not supported.

- 65 -

CREG<4>	Reserved	This is a reserved bit. It should always be set to a value of 1. A value of zero in this field is not supported.
CREG<5>	Reserved	This is a reserved bit. It should always be set to a value of one. A value of zero in this field is not supported.
CREG<6>	Reserved	This is a reserved bit. It should always be set to a value of one. A value of zero in this field is not supported.
CREG<7>	SEQ/486#	Setting this bit low implies that 486 burst-order is selected for data transfers between system memory and cache. Otherwise, sequential burst order is assumed.

5 Expansion Register (Index Address = 01) - This register defines the physical operation address space of each controller 70 in cascade mode. Each controller 70 will provide control for 64K byte of cache memory. Provisions in XREG allow an expansion configuration of up to 256K byte of cache memory using four controllers.

Bit Position	Name	Definition
XREG<0>	INV	Setting this bit will cause an invalidation of all cache data entries. The controller 70 will write back all dirty data associated with each entry before invalidation. At the end of the write back operations, the controller 70 will clear this bit.
XREG<2:1>	CSIZE<1:0>	These two bits identify the cache memory size as follows:
CSIZE<1:0>		Cache Size
00		64K byte
01		128K byte
10		256K byte

- 66 -

XREG<3>	DIS	Setting this bit to one will disable the controller 70. All bus cycles will be treated as NCA (Non-Cacheable Address) system cycles. If this bit is used to dynamically disable the cache during normal operation, a cache flush should proceed the disabling, in order to flush any dirty data in the cache.
XREG<5:4>	TCF<1:0>	These two bits define the tag array associativity configuration as follows: TCF<1:0> Assoc. Tag Add. Org. 00 DIRECT MAPPED 1X2048X12 01 2-way associ- 2X1024X13 ative
XREG<7:6>	Reserved	These bits are reserved.

Protection Register (Index Value = 10 Hex)- The
5 protection register defines the operation of the four
available protected regions. Each protected region is
associated with an NCA bit and a CWP bit. Note that
either the NCA bit or the CWP bit may be set for a
protected region, but not both. The bit definition and
10 nomenclature are listed as follows:

Bit Position	Name	Definition
PREG<7,5,3,1>	NCA<4:1>	The setting of each of these four bits defines one of the four protected regions of the cache controller to be non-cacheable. NCA<1> high will disable cacheing in the address range defined by PR1S and PR1E. Likewise, NCA<2>, NCA<3> and NCA<4> will disable cacheing in regions two, three and four.

- 67 -

PREG<6,4,2,0> CWP<4:1> Setting this bit defines the corresponding region as cacheable, but write-protected. Writes to an address region with the CWP bit set will be bypassed to the system by the controller 70. By use of this bit, system and video BIOS may be safely cached.

Protected Address Region Registers (Index Value = 2x to 3x Hex) - As previously discussed, the controller 70 can protect up to four address regions. Each protected region is defined through two 16-bit registers and two operating mode bits, NCA and CWP. The starting (low) address register and the ending (high) address register identify the address range of the protected regions. The starting and ending addresses of all four protected regions may be defined to 4K byte boundaries. For example, the PR1S and PR1E registers define the start and end of the first protected region. The second through fourth regions are similarly defined. These addresses are exclusive, and therefore not intuitive. For example, to use region 4 as a protected region from address 40000 hex to 7FFFF hex, the start address (PR4S) should be loaded as 003F hex and the end address (PR4E) should be loaded as 0080 hex.

The physical address of each bus cycle is compared against the protected address ranges. An address within any of the range will modify controller response according to the values of the corresponding NCA and CWP bits as follows:

	NCA	CWP	Meaning
25	0	0	Normal (cacheable) region
	0	1	Cacheable write protected region
	1	0	Non-cacheable address region
	1	1	Undefined state; do not use

- 68 -

An address not contained in any of the four protected regions is assumed to be a cacheable address. Settings such that the value of the starting address is larger than that of the ending address should be avoided;

5 indeterminate effects may result. In addition, a region defined with equal starting and ending address will be cacheable, regardless of the NCA and CWP bits. Note that it is possible to define overlapping protected regions, each with different mode definitions. In this case, the

10 priority will be as follows:

1. CWP (Cacheable Write Protected) regions
2. NCA (Non-Cacheable Address) regions
3. Normal cacheable regions

Default Setting for Control Registers at Reset -

15 After reset, the cache controller 70 is essentially ready for use in a PC environment. For easy system integration and best performance, the default register values prepare the controller for use in write-back mode using 486 CPU burst order to memory and assuming use of cache memory 72.

20 For proper operation, all that is necessary for the BIOS or operating system is to set bit 0 of the control register, in order to enable the cache. Three of the protected regions are used at reset, leaving the fourth available for use. The protected address regions and

25 their corresponding protection register bits are initialized as follows:

Register		Value <7:0>	
30	Control Register	1111	1101
	Expansion Register	0001	1000
	Protection Register	0001	0110

Register	Range	Value A<27:12>	Mode NCA:CWP	Reason
1st region start	640K	009F	1:0	top
1st region stop	1 Meg	0100	Non- cacheable	384K of 1 Mb

- 69 -

	2nd region start	768K	00BF	0:1	VGA
	2nd region stop	800K	00C8	Cacheable	BIOS
				Write	
				Protect	
	3rd region start	960K	00EF	0:1	System
	3rd region stop	1 Meg	0100	Cacheable	BIOS
				Write	
				Protect	
5	4th region start	N/A	0000	0:0	User-
	4th region stop	N/A	0000	Cacheable	defin-
					able

Before loading any of the non-cacheable region support registers, the cache should be temporarily disabled. This can be done by clearing the CE bit in the control

10 register. After non-cacheable regions are changed, the cache should then be invalidated by setting the inv bit in the control register. This avoids any data coherency problem as a result of the non-cacheable region redefinition. After the invalidation process is complete,

15 the cache can safely be re-enabled.

FUNCTIONAL DESCRIPTION

This section discusses how the controller 70 responds to the different types of bus cycles generated by the 486 CPU. The discussion covers normal cacheable memory

20 reference (read/write) operations, locked, interrupt acknowledge and halt/shutdown cycles. Special 486 CPU cycles like Flush and Write-back cycles for supporting the 486 CPU INVD and WBINVD instructions are also described. The interface with cache memory 72 burst SRAMs is also

25 discussed.

The controller 70 supports 486 CPU systems with the cache memory 72 Burst-RAMs. The controller 70 only supports write-back mode for 486 systems. Write-through is supported on a cycle to cycle basis through the PWT

30 input. The controller 70/cache memory 72 support burst reads for 486 cache line fills. The controller 70 will follow 486 style address sequence on the 486 CPU local bus. On the system bus, either 486-style or sequential

- 70 -

address sequence is supported. Use of cache memory 72 allows miss operations and write-back cycles to be carried out in parallel with hits.

- 486 CPU Bus Cycle Definitions - M/IO#, D/C# and W/R#
- 5 are the primary bus cycle definition signals from the 486 CPU. These signals are driven valid in T1, as ADS# is asserted. M/IO# distinguishes between memory and I/O cycles. D/C# distinguishes between data and code cycles. W/R# distinguishes between write and read cycles.
- 10 Three other 486 signals provide cycle definition to the controller 70. These signals are:
1. The PCD (Page Cache Disable) pin, which reflects the enabling of the 486 CPU internal cache for the current cycle. The great majority of
15 operations will occur with the 486 CPU internal cache turned on. This is indicated by the 486 PCD output being de-asserted (low). PCD will be asserted high by the 486 if its internal cache has been turned off by software, either entirely
20 or on a per-page basis.
 2. The PWT (Page Write Through) pin, which indicates to the controller 70 that the corresponding write cycle should be directed to the system, as well as updating the contents of
25 the data cache if a hit occurs. Like the PCD pin, the PWT pin is used by 486 CPU software on a per-page basis.
 3. The LOCK# pin, which asserted indicates that the 486 CPU is performing a read-modify-write
30 operation over several bus cycles. The 486 CPU should retain ownership of the bus while LOCK# is asserted.

Table IV shows the encodings for the various bus cycles that occur in 486 CPU systems. Halt cycles have
35 been moved to location 001 from location 101 for the 386 CPU. Location 101 is now reserved in the 80486.

- 71 -

Table IV. 80486 Bus Cycle Definitions

	M/IO#	D/C#	W/R#	80486 Cycles
	0	0	0	Int Acknowledge
	0	0	1	Special Cycles
5	0	1	0	I/O Read
	0	1	1	I/O Write
	1	0	0	Memory Code Read
	1	0	1	Reserved
	1	1	0	Memory Data Read
10	1	1	1	Memory Data Write

486 Special Cycle Definitions - In addition to the M/IO#, D/C# and W/R# signals, the 486 CPU provides four special bus cycles to indicate that certain conditions have occurred internally or certain instructions have been executed. These four special bus cycles are defined by the byte enable signals when M/IO# = 0, D/C# = 0, and W/R# = 1. Table V shows the encodings of the 486 special bus cycles.

Table V. 80486 Special Bus Cycle Definitions

20	BE3#	BE2#	BE1#	BE0#	Special Cycles
	1	1	1	0	Shutdown
	1	1	0	0	Flush
	1	0	1	1	Halt
	0	1	1	1	Write Back

Differences Between the controller 70 and i486 CPU - The controller 70 presents a very 486 CPU-like interface to system logic. The great majority of system interface pins have the same name and functionality as their 486 CPU counterparts. Because of these features, designing the controller 70 into a system is very straightforward. Glue logic to design in the chipset is minimal. In addition, hooks can be used to drive optional address transceivers and latches on the system side. These devices may be used if additional drive capability is desired; however, the controller 70 system side specifications assume 100 pF of loading. The SA OE# and SA DIR control signals are No Connect pins if address transceivers are not used.

- 72 -

The controller 70 supports all i486 CPU functionality on the host (local CPU) side. A high-speed 32-bit interface is used, with the great majority of cycles completing in two clock cycles. On the system side, the controller 70 has an identical list of pins as does the i486 CPU, except for the following pins:

- PCD - Page Cache Disable
- PWT - Page Write-Through
- AHOLD - Address Hold
- 10 BS8# - Bus Size 8
- BS16# - Bus Size 16
- PLOCK# - Pseudo Lock
- KEN# - Cache Enable

These signals are not included in the system-side controller 70 architecture for various reasons:

The 486 CPU architecture intends the PCD and PWT signals to be used by an external cache, and the system bus has no need of these signals.

AHOLD is not needed in current system design. The preferred method of invalidating i486 CPU cache lines is through the SHOLD/SHLDA protocol, and will be discussed later.

BS8# and BS16# are not supported; systems must interface to the controller 70 with a 32-bit interface.

25 PLOCK# is rarely used in systems. If it is desired to include PLOCK# in a system, a fast AND gate may be used to connect the LOCK# and PLOCK# outputs of the i486 CPU. The resulting AND will then be used as the LOCK# input of the controller 70.

30 KEN# is not necessary on the system side, as the controller 70 contains register support for four protected address regions, all of which may be either entirely non-cacheable or read cacheable write-protected. Addresses will be decoded by these registers and KEN# returned to 35 the i486 CPU in T1, in order to avoid any performance degradation. Use of the controller 70 in a system eliminates the need for any high-speed address decode

- 73 -

PALs. In addition, the programmable register approach for non-cacheable regions is more flexible than PAL implementations, since the end-user can dynamically tailor cacheability.

- 5 BOFF#, or Back Off, is not included in the system-side controller 70 architecture. This functionality will be included in the next generation of the controller 70 family of controllers.

In addition, the controller 70 has a bus snooping
10 feature and the ability to intervene on system snoop reads, when the requested data is both present in the cache and also 'dirty'. As a result, the controller 70 has three additional pins to support this functionality:

- SNPBUSY - Snoop Busy
15 SMEMWR - System MEMory Write Read
SMEMDIS - System MEMory DISable

The functionality of these pins will be discussed below.

Table VI

20 Summary of controller 70 Response to CPU Cycles

System Cycles:	NCA, I/O, Halt/Shutdown, INTA cycles and controller 70 disabled No cache update occurs controller 70 re-drives cycle on system bus one clock later SBLAST# is driven low, regardless of state of BLAST# System returns SRDYI# or SBRDYI# to terminate; passed to 486 CPU as RDYO# No write buffering, except NCA writes will be buffered KEN# is deasserted to CPU; no line fill will occur
----------------	---

- 74 -

Local Bus Cycles: Controller 70 Control Register
Read/Writes, Weitek Cycles, and CWP
Write Cycles
No cache operations
For control register read/writes,
controller 70 returns RDYO#
For Weitek cycles, Weitek returns
RDY#; passed to 486 CPU as
RDYO#
CWP Write Cycles are terminated
in two clocks with no system
cycle

Write-Through
Cycles: PCD/PWT/Locked Write Cycles
Read Hit/Miss: N/A
Write Hit: Update CPU write data
into cache data
array
controller 70 returns
RDYO#
to CPU in zero wait
states
Buffered write con-
tinues on system
bus until SBRDYI#
or SRDYI#
Write Miss: No update of CPU
write data into
cache data array
controller 70 returns
RDYO#
to CPU in one wait
state
Buffered write
continues on system
bus until SBRDYI#
or SRDYI#

- 75 -

Normal Cycles:

Normal Read/Write Cycles and
PCD/PWT/CWP/Locked ReadRead Hit: controller 70/cache
memory 72's supply
data in zero wait
states
controller 70 returns
BRDYO#to CPU; Host port
transfer terminates
with BLAST#CPU assertion of
BLAST# terminates
host port transferFor normal/PWT reads,
KEN# asserted to
cause i486 CPU
cache line fillKEN# deasserted for
PCD and CWP reads;
no line fill will
occurRead Miss: System Quad Fetch;
data bypassed to
CPU, also latched
in memory update

register set 116

System returns 4

SRDYI# or 4

SBRDYI#; either
passed to CPU as
BRDYO#controller 70 asserts
SBLAST#on fourth (last)
cycleFor normal/PWT reads,
KEN# asserted to
cause i486 CPU
cache line fillKEN# deasserted for
PCD and CWP reads;
no line fill will
occur

- 76 -

CPU assertion of
 BLAST# terminates
 host port transfer
 controller 70 assertion
 of
 QWR signal writes
 memory update register
 set 116 into cache
 data array
 Tag miss requires any
 dirty data in write
 back register set 118
 to be written back
 to system
 Write Hit: controller 70 returns
 RDYO#
 in zero wait states
 (two clocks)
 CPU write data is
 written into cache
 data array
 no system cycle occurs.
 Write Miss: Line to be replaced
 is stored into
 write back register
 set 118
 CPU write data is
 written into cache
 data array
 System Quad Fetch;
 data not bypassed
 to CPU, but latched
 in memory update
 register set 116
 System returns either
 4 SRDYI# or 4
 SBRDYI#
 controller 70 asserts
 SBLAST#
 on fourth (last)
 cycle
 controller 70 assertion
 of
 QWR signal writes
 memory update register
 set 116 into cache
 data array
 Tag miss requires any
 dirty data in write
 back register set 118
 to be written back

Controller 70 Response to 486 CPU Cycles - As shown
 in Table VI, the controller 70 distinguishes between four

- 77 -

main classes of cycles. These four classes are system cycles, local bus cycles, write-through cycles, and normal (cacheable) cycles.

I) System Cycles: The controller 70 detects a bus cycle as a system cycle through either one of the following conditions:

1. An NCA cycle. An NCA cycle is a read/write cycle to an address defined in the protected address region registers with the NCA (non-cacheable address) bit set.
2. A cycle which is an I/O read or write, an interrupt acknowledge, or halt/shutdown cycle.
3. Controller 70 disabled. If the CE bit in the control register is cleared (0), all cycles will become system cycles.

For this class of cycles, the controller 70 forwards the address and bus cycle control signals to the system bus without performing a cache access. All read cycles are treated as read misses except that the cache directory and the cache data array are not affected. All writes are treated as write misses. NCA write cycles will be buffered. For all system cycles, KEN# is de-asserted to the 486 CPU. This will prevent the returned data from being cached in the CPU.

SBLAST# will be low in ST1, regardless of the state of BLAST# from the 486 CPU. For most system cycles, this will have no impact, as the 486 CPU BLAST# output will be low for I/O, INTA, Halt/Shutdown, and normal write cycles. However, the effect of driving SBLAST# low means that non-cacheable address reads cannot be burst from system memory. By defining an address range to be NCA, it is also defined to be non-burstable.

Figure 34 shows a system read cycle, where the cycle definition is passed on to the system. The controller 70 asserts SADS# the clock after ADS# was asserted. The Bypass signal is asserted to allow returned system data to be passed back to the 486 CPU in the same clock. HPOEx#

- 78 -

is enabled to allow the read data to be passed back to the 486 CPU.

Figure 35 shows an I/O write cycle, which is passed on to the system without being buffered. SRDYI# and
5 SBRDYI# are passed back to the 486 CPU as RDYO# to terminate this cycle. In the ST2 state, SPOE# is asserted low to enable the system port. MALE is asserted high by the controller 70 in ST2 to inhibit the cache memory 72 write operation.

10 I/O Cycles - I/O cycles are passed on to the system bus, and terminated when the system asserts SRDYI# or SBRDYI#. Either of these signals are passed to the CPU as RDYO#. I/O cycles are not buffered. I/O cycles will not produce any cache operations.

15 INTA (Interrupt Acknowledge) Cycles - The 486 CPU generates interrupt acknowledge cycles in locked pairs. The controller 70 will re-drive these to the system, with the same encoding as on the 486 CPU. Also like the 486 CPU, the state of A2 will allow system logic to
20 differentiate between the two INTA cycles. A2 will be driven high during the first INTA cycle, and low for the second. SLOCK# will be asserted between and during both of these cycles.

Data parity for the two interrupt acknowledge cycles
25 is not checked by the 486. External hardware must acknowledge each interrupt acknowledge cycle through SRDYI# or SBRDYI# assertion. The controller 70 will invoke the Bypass signal and HPOEx# to the cache memory 72 during the second INTA cycle so that interrupt vectors are
30 passed from the system bus to the local processor bus.

Halt/Shutdown Cycles - The controller 70 treats halt/shutdown cycles as system cycles. During
halt/shutdown cycles, the controller 70 duplicates the encoding of the host processor on the system memory bus.
35 External hardware should acknowledge halt/shutdown cycles

- 79 -

through SRDYI# or SBRDYI# assertion. During halt/shutdown cycles, the controller 70 will recognize SHOLD from the system memory bus and will respond by floating the system address, system control definition signals and asserting
5 SHLDA.

NCA (Non-Cacheable Address) Cycles - The address of each bus cycle is compared against the contents of the protected address region registers to determine cache-ability of the cycle.

10 If the address of a 486 CPU cycle is within a protected address region with the NCA bit set, the cycle is determined to be an NCA cycle. The cycle will be forwarded to the system bus without any cache operation taking place. No data will be supplied from the cache in
15 this case. In addition, KEN# is asserted high to the CPU to prevent data on read cycles returned from the system from being cached. KEN# being returned high in T1 prevents the 486 CPU from performing any throwaway line fill cycles, which would degrade performance.

20 Unlike the other cycles which make up the system class, write cycles to non-cacheable addresses are terminated in two clocks by the controller 70 by assertion of RDY0# to the 486 CPU. These buffered writes are completed on the system bus when either SRDYI# or SBRDYI#
25 is returned by system logic. The Bypass signal will be asserted in either the NCA read or write cycles until the cycle is terminated. For NCA reads, the activation of Bypass allows returned system data to pass to the CPU. For NCA writes, Bypass allows the buffered write data
30 contained in write register 120 to be driven to the system.

Figure 36 details an NCA write cycle. The write cycle is buffered and terminated on the local CPU bus in two clocks by the controller 70. Note that like a normal
35 cache write hit, HPWEx# is asserted low in ST2. However, the write to the cache data array is inhibited by

- 80 -

assertion of MALE. The write data is latched into write register 120. Write register 120 will act as the write buffer in this cycle. The cycle continues on the system side until terminated by SRDYI# or SBRDYI#. The Bypass
5 signal stays activated until SRDYI# or SBRDYI# is asserted. Continued assertion of the Bypass signal after the cycle has been terminated on the local CPU side indicates that the data in write register 120 is being driven to the system bus.

10 II) Local Bus Cycles: - The controller's 70 second class of cycles are Local Bus cycles. Local bus cycles are not passed on to the system, nor do they cause a cache hit/miss determination. Local bus cycles consist of reads
15 bus cycles, and writes to cacheable write-protected regions.

Controller 70 Register Reads/Writes - Control words are read and written to the registers of the controller 70 through a two-step index addressing process. First, a
20 write cycle is performed to the address of the controller 70. The data for this write cycle is the index address of the desired register. A2 should be low for this cycle. This indicates to the controller 70 which register is to be read or written to.

25 A second cycle is then performed to the controller 70. This cycle performs the actual read or write to the control register. The data lines for this cycle contain or return appropriate read/write data. A2 should be high for this cycle.

30 Figures 37 and 38 show reading and writing to one of the control registers of the controller 70. For these cycles, MCCSEL# should be decoded and asserted to the controller 70 in the T1 state. The controller 70 will return RDYO# to terminate these cycles. Both reads and
35 writes will take on wait state.

- 81 -

Weitek Bus Cycles - The host CPU may execute bus cycles that access the Weitek math coprocessor (4167) in some 486 CPU systems. The controller 70 simply ignores these accesses and does not initiate any activity in the cache or on the system bus. The controller 70 recognizes bus cycles intended for the 4167 through the pattern of A<31:25> being <1100000>. The 4167 acts as a device in a reserved memory space. The 4167 generates its RDY# signal to the controller 70, which is passed onto the 486 CPU.

Detection of a Weitek coprocessor cycle will unconditionally deassert the KEN# input to the 486 CPU. Data to and from the Weitek coprocessor is regarded as non-cacheable.

If desired, the local RDY# input from the Weitek coprocessor may be operated in an asynchronous mode. In this mode, RDY# is latched and sent to the 486 CPU in the following clock, instead of in the same clock. This mode should be used if the synchronous mode RDY# setup time cannot be otherwise met. Bit 1 in the controller 70 control register controls this mode.

CWP Write Cycles - Write cycles to cacheable write-protected regions will be terminated by the controller 70. RDY# will be returned to the i486 CPU to complete the cycle in zero wait states, and no system cycle will occur. This will relieve system designers from the duty of decoding and inhibiting writes to regions of memory where code has been shadowed from EPROM to DRAMs.

III) Write-Through Cycles: - The controller 70 defines a third class of cycles to be the write-through cycles. The controller 70 detects a write-through cycle through the following condition:

1. The 486 CPU asserts either the PCD, PWT, or LOCK# pin for a write cycle.

Note that read cycles, by definition, never qualify

- 82 -

as write-through cycles. PCD, PWT, and Locked reads are treated as normal read cycles.

When a write-through cycle occurs, the controller 70 will always generate a write to the system bus, regardless 5 if the cycle is determined to be a cache hit or miss.

Write-through cycles which are hits will generate updates to the cache RAM array section 100. In addition, the controller 70 will buffer the write. RDYO# will be returned to the 486 CPU in the first T2 to terminate the 10 write on the CPU bus. The cycle will continue on the system side until the assertion of SRDYI# or SBRDYI#, which indicates that the system has accepted the write data.

Write-through miss cycles will also be buffered 15 writes, although RDYO# will be asserted after three clocks, instead of two as for write hits. Unlike normal write misses, write through miss cycles will neither update the cache nor generate system quad fetches.

IV) Normal (Cacheable) Cycles: - The fourth class of 20 cycles are the normal cacheable cycles. These cycles will be the great majority of cycles which occur. Normal cacheable cycles are the default cycles, and are assumed if a cycle does not fit into any of the previously described categories. Specifically, the controller 70 25 detects a normal cacheable cycle under either one of the following conditions:

1. A read cycle which is not an I/O, Interrupt Acknowledge, Halt, or Shutdown cycle. In addition, reads whose addresses are contained in 30 the protected address region registers without the NCA bit set are cacheable cycles.
2. A write cycle which is not an I/O, Interrupt Acknowledge, Halt/Shutdown, PCD, PWT, or Locked cycle. In addition, writes whose addresses are 35 contained in the protected address region registers without either the CWP or NCA bits sets are cacheable cycles.

- 83 -

READ OPERATIONS

Cacheable Read Hit Operations - When the 486 CPU initiates a memory data/code read cycle, the assertion of ADS# triggers a compare cycle in the tag array of the controller 70 if the cycle is of the type that can be cached. If a match exists between the cycle address and one of the tag directory entries, a hit results.

A read hit will result in either a single, double, or quad transfer to the 486 CPU. The number of transfers depends on the state of the internal 486 CPU cache, as reflected by the PCD pin. The meanings of the two states of PCD are described below:

1. **PCD=0:** If the 486 CPU internal cache is enabled and a hit occurs, a quad doubleword transfer will result, in order to fill a line of the 486 CPU. KEN# is returned active (low) twice to the 486 CPU, first in T1 to initiate the line fill, and also the clock before the final transfer. The quad transfer will take five clocks with cache memory 72 burst-RAMs. The demand doubleword will be returned first followed by the other three remaining doublewords in the same line, following the 486 address order. BLAST# will be asserted by the 486 CPU during the fourth transfer. A quad burst transfer from 82C443 Burst-RAMs is shown in Figure 39).
2. **PCD=1:** PCD being asserted high by the 486 CPU indicates no line fills will occur in the 486 CPU cache. KEN# will be returned high to the CPU in the T1 state. Either a single, double, or quad transfer will occur in this case, depending on the BLAST# output of the

- 84 -

5 CPU. The controller 70 will monitor BLAST# from the 486 as each transfer is completed. BLAST# assertion (low) indicates that the cycle associated with the corresponding BRDY0# is the last data cycle. BLAST# assertion will terminate the transfers. A single transfer is shown in Figure 40.

The first demand word will require two bus states.
10 For each of the remaining doublewords, only one bus state is required. BRDY0# will be asserted for each returned doubleword. Single doubleword transfers will require two clocks and double transfers three clocks.

A read cycle to a CWP region will result in the
15 controller 70 returning the KEN# output deasserted (high) to the i486 CPU. This will prevent the corresponding data from being cached inside the i486 CPU internal cache.

The details of hit transfers are as follows. The controller 70 will assert CCSx#. The bank being read is
20 enabled by assertion of HPOEx#. The controller 70 will return BRDY0# to terminate each transfer to the 486 CPU, until BLAST# signals the end of the cycle. The first access will come from the cache data array. However, read hold register set 114 will be loaded with the data to
25 satisfy the entire quad read. Data for the remaining second through fourth transfers will be read from the read hold register set 114. To connect the output of read hold register set 114 to the host port outputs, signal RHSTB will be asserted high in the second ST2 state. For each
30 transfer, the controller 70 will drive HA<3:2> to valid levels to provide 486-style address sequencing. HA3 and HA2 select a doubleword from the four in each line.
Although the 486 CPU provides identical functionality as HA3 and HA2 with its A3 and A2 signals, HA<3:2> should be
35 used, since their valid delays are much shorter than those of A<3:2> from the 486 CPU.

- 85 -

Cacheable Read Miss Operations - The tag comparison for a cycle may indicate the occurrence of a read miss. There are two kinds of misses, a tag miss and a line miss. A tag miss results when the tag lookup does not produce a
5 match, or the tag valid bit is not set. A line miss occurs if the tag lookup produces a match and the tag valid bit is set, but any of the four doublewords associated with the line are not valid. Tag misses and line misses are treated differently by the controller 70.

10 Cacheable read misses generate a system quad fetch. On system quad fetches due to tag misses, data that is retrieved from the system may replace older data existing in the cache. If any of the data which is replaced is both valid and dirty, one or two write-back cycles will
15 occur.

On line misses, there is a tag match, but the line is not wholly existent in the cache data array. Data received from the system will be merged with valid dirty data already contained in the line. As no dirty data is
20 replaced, no write-back cycle(s) will occur from line misses.

The controller 70 examines the LRU bits of the target entries to select which of the two banks is to be replaced. If the "tag-valid" bit for the entry chosen to
25 be replaced is set, the controller then checks if any one of the eight doublewords of the two lines which correspond to the selected tag entry are set "valid" and "dirty". Any such words within either line means that the corresponding line will have to be written back to main
30 memory by later write-back cycle(s). All valid bits as well as the "dirty" (altered) bit for the lines will then be reset.

The controller 70 will latch the read miss address into the cache memory 72's miss address register 110. The
35 miss address together with the LRU bits select the data line to be replaced. This data is then latched into the write back register set 118 of the cache memory 72.

- 86 -

On the host side, the 486 PCD pin affects read miss operations similarly to the read hit case. PCD being low indicates a cache line fill for the 486 and cache memory 72 will occur. The controller will assert KEN# twice to
5 the 486 CPU to perform a line fill, before the first and fourth transfers are completed. If PCD is high, no line will occur in the 486 CPU cache, although a line fill will occur in the cache memory 72. KEN# will be returned high to the 486 CPU in this case.

10 Cacheable read misses initiate quad fetches on the system side, in order to bring four doublewords into the cache data RAM and fill a cache line. This quad fetch will continue, eve if the 486 CPU does not require all four doublewords (i.e., the 486 CPU internal cache is
15 turned off). The CPU may terminate the fetch on the host side after one or two transfers, while the quad fetch from the system continues until completion. The quad fetch is finished in order to increase the cache hit rate and, as a result, overall performance. The controller 70 will
20 assert SBLAST# on the fourth transfer, to indicate the completion of the cycle.

Figure 41 shows a read line miss with PCD asserted by the 486 CPU. A system quad fetch results and completes, although the 486 CPU terminates the cycle on its local bus
25 after only two transfers. Note that another read or write hit could then be processed on the local bus while the system quad fetch completes. The write-back architecture of the controller 70/cache memory 72 isolate local CPU and system bus processing.

30 In Figure 41, signal MALE is asserted high in T2 to indicate that a miss has occurred. The assertion of MALE latches the read miss information. Bypass is asserted high to allow data from the system quad fetch to pass on to the 486 CPU in the same clock. The addition of the
35 Bypass propagation delay means that valid data setup time for the cache memory 72 must be greater than that of the 486 CPU. HPOEx# goes low to enable the host port outputs

- 87 -

to bypass data received from the system.

During a system quad fetch, system logic can assert either SRDYI# four times or SBRDYI# four times to complete the transfer. A combination of SRDYI# and SBRDYI# 5 assertions (i.e., interrupted burst) is not supported by the controller 70. To sustain high performance, the controller 70 will pass either ready input from the system to the 486 CPU as BRDY0#. BLAST# is monitored as each doubleword is transferred. BLAST# assertion will 10 terminate the cycle on the local CPU side.

The system quad fetch from main memory will not be written directly to the cache data array. Instead, fetched data will be loaded into memory update register set 116. Since the line may have been partially valid and 15 contained some dirty doublewords which should not be overwritten, incoming doublewords from system memory will be qualified before they can be written into memory update register set 116. The controller 70 will assert the DW# signal to indicate to the cache memory 72 that the fetched 20 doubleword may be safely written into memory update register set 116. Inactivation of DW# inhibits dirty and valid miss data from being overwritten by system quad fetch data.

There is a "mask" bit associated with each of the 25 four words of the memory update register set 116. The activation of this "mask" bit disables the corresponding memory update register set 116 word from updating the cache memory 72 RAM array section 100. The update clears all "mask" bits. At the completion of the system quad 30 fetch, the controller 70 will write the contents of memory update register set 116 into the cache memory 72 data array by asserting the QWR signal. Associated "valid" bits of the corresponding line will be set.

The controller 70 will send the appropriate control 35 signals to turn the data transceiver to receive mode. The address sequence to main memory will either follow 486 style address sequence or sequential address sequence.

- 88 -

If the burst order at main memory interface is the same as the 486 (no re-ordering is occurring), the Bypass signal will be activated and the incoming data from main memory will be sent to the local bus through the bypass
5 path inside the cache memory 72. At the same time, the data will be latched in a cache memory 72 register, to be updated into the cache memory 72 once the entire line is received.

Figure 42 details a read line miss, with the
10 controller 70 responding by activating a quad fetch on the system memory bus. Assertion of SADS# and the other control signals initiate the quad fetch. Bypass is activated, so that the data being returned from the system is forwarded on to the 486 CPU in the same clock. Because
15 the read data must propagate through the cache memory 72 Bypass path and still meet 486 CPU setup time specification, the read data setup time to the cache memory 72 is longer than that of a 486 CPU by itself.

To assist the system in performing one clock bursts,
20 the SA3 and SA2 (system address 3 and 2) signals become valid early in each T2 state, much sooner than the 486 CPU would provide them. In addition, the read data is latched into the cache memory 72 memory update register set 116 register. At the completion of the transfer, the
25 controller 70 activates SBLAST# to terminate the cycle. The data in memory update register set 116 is written into the cache memory 72 RAM array section 100 by activation of the QWR (Quad WRite) signal.

Figure 42 also details a complexity that will
30 sometimes occur. On read line misses, the line is partially present in the cache memory 72 data array, and some of these doublewords may be dirty. When the corresponding address for dirty doublewords is generated during the system quad fetch, the data should be returned
35 from the cache data array instead of the system. For these dirty doublewords, the Bypass signal will be dynamically turned off (low) and the dirty data will be

- 89 -

correctly supplied to the 486 CPU from cache memory 72. This process is transparent to the system, which should supply data for all transfers of a quad fetch.

This occurrence is shown during the third transfer in Figure 42, where Bypass is de-asserted and the data that the 486 CPU receives comes from the cache memory 72 read hold register set 114 instead of the system. Note that to be prepared for this case, signal RHSTB is asserted high in T2 to load read hold register set 114 with the cache data for the miss address line.

Figure 43 details a read line miss with reordering. A difference in address order between the main memory interface and the 486 will result in the first demand word to be sent to the local bus through the cache memory 72 bypass path. The second and third doublewords within the line will be latched into cache memory 72 memory update register set 116 as well as being sent to the 486 CPU. However, since reordering is being done, RDYO# will not be asserted to the 486 CPU and it will not latch these doublewords. The fourth doubleword is then returned from main memory, passed to the CPU by way of the cache memory 72 bypass path and latched into memory update register set 116 simultaneously.

The subsequent two remaining doublewords (second and third doublewords transferred) now in memory update register set 116 with 486 burst-order. The controller 70 will generate BRDYO# for these two remaining returned data cycles. Note that as in the previous figure, part of the line may be valid and dirty. The controller 70 will de-assert Bypass and supply the 486 CPU with the correct dirty data during the appropriate transfers.

Write-Back Cycles Due to Read Tag Misses: - The new data to be brought into the cache from the system during read misses replaces older data residing in the cache. However, the write-back architecture of the controller 70 requires additional considerations when this old cache

- 90 -

data is replaced. A direct replacement algorithm would cause data to be lost if the replaced data is valid and dirty. Data is marked dirty if it has been modified by the 486 CPU but not yet been copied back to main memory.

5 Because of this, read tag misses and the subsequent loading of a new line of data into the cache from the resulting system quad fetch may be followed by one or two write-back cycles to main memory at the end of the quad fetch. For high performance, the controller 70 allows
10 these quad writes to be burst to main memory, if the system memory controller is capable of receiving such bursts.

The controller 70 will generate either zero, one or two quad write replacement cycles to main memory, which
15 may be burst by the system. No write replacement cycle will be generated if the selected tag entry is invalid or both lines associated with the selected tag entry contain no dirty data.

If any of the doublewords for a given replaced line
20 are marked dirty and valid, a write-back cycle will occur and write all four doublewords of the line to main memory. Quad fetches due to misses and write-back replacement cycles will not abort due to subsequent miss cycles. Address order on quad-writes may follow either 486 burst
25 order or sequential burst order. Unlike write-back cycles which occur due to flush operations, the CALE signal will not be generated for write-backs due to misses, as the needed address is already contained in the cache memory 72 miss address register 110.

30 Quad write transfers will contain a mixture of dirty and non-dirty doublewords. As a performance enhancement, the controller 70 provides a signal to avoid performing writes which would contain non-dirty data (which the system already contains). The DW# signal is driven valid
35 in all T2 states of write cycles. To simplify system design, DW# will become valid early in each state.

DW# should be incorporated into two parts of system

- 91 -

memory logic. First, DW# should be part of the write enable logic. DW# being low for a transfer indicates that the corresponding data is indeed dirty and should be written into the system memory array. Writes in which DW# 5 is driven low are processed normally by the system.

However, when the DW# is driven high, the data appearing on the cache memory 72 data outputs is not dirty. The system may, at its option, accept this write data.

10 To add additional system performance, DW# can be incorporated into the SRDYI#/SBRDYI# logic in order to quickly terminate non-dirty writes. The state of DW# being high can be used in combinational logic to terminate non-burst dummy write transfers in two clocks and burst 15 dummy transfers in one clock. Because DW# becomes valid early in each T2 bus state, it can easily be incorporated into system logic without being the most critical path.

At the end of the first write-back cycle, if any of the four double-words from the other line associated with 20 the replaced tag are marked valid and dirty, the controller 70 will invoke another replacement burst write-back cycle through MWB activation. WB STB will be asserted. MWB assertion allows the cache memory 72 internal miss address register (miss address register 110) 25 counter to be pointed at the desired replace line and the assertion of WB STB allows the next accessed replacement data to be latched into write back register set 118. This data will later be written back to main memory through a second write-back cycle.

30 System assertion of either SRDYI# four times or SBRDYI# four times to terminate the transfers is allowed. Combinations of SRDYI# and SBRDYI# to terminate the transfers of a quad write are not supported. Once the system has asserted either SRDYI# or SBRDYI# to complete 35 the first transfer, it must assert the same signal three more times to finish the transfer. The controller 70 will assert SBLAST# during the fourth transfer to indicate

- 92 -

completion of the quad write.

Figure 44 details a write-back cycle to memory, where both lines associates with a tag entry contain valid and dirty data. On the system side, SADS# and all system
5 cycle definition signals are driven valid in T1 (with the exception of DW#, which becomes valid in T2). Note DW# toggling each transfer to indicate dirty/non-dirty status. SPOE# is asserted in ST2 to enable the system port 112 of the cache memory 72.

10 WRITE OPERATIONS

Write Hits

For best performance, the controller 70 supports 486 CPU systems only with write-back mode. Write-through mode is supported on a cycle-to-cycle basis by monitoring the
15 PWT signal on the 486. In write-back mode, the controller 70 updates the cache memory without updating the system memory if a cache hit occurs. Main memory is updated only when a dirty line is replaced.

Figure 45 shows a cacheable write hit operation. The
20 cycle is terminated in two clocks by assertion of RDY0# to the 486 CPU. The write data is simply latched into the cache data array by assertion of HPWEx# in T2. The associated valid and dirty bits for the double word are set in the controller 70 tag array. BLAST# will always be
25 asserted on the first transfer, as only scalar writes are supported by the 486 CPU. The 486 CPU cannot burst write more than 32 bits.

Write Miss Operations

Cache memory 72 write-back architecture filters out
30 most of the write traffic induced by the 486 CPU internal cache write-through policy. System write cycle latency is eliminated as the performance bottleneck. Furthermore, use of cache memory 72 allows write misses to be followed by zero wait states read/write hits with no idle bus
35 clocks.

Write Miss Cycles

- 93 -

In order to achieve maximum performance, the controller 70 will essentially treat write miss cycles as write hit cycles, by writing the data directly into the cache and replacing one line in the cache. A one-clock latency is required to first move the data being replaced from the cache data array into the write back register set 118, for later write-back if dirty data was present in the replaced line. As a result of this one-clock latency, write miss cycles will occur in three clocks, instead of two.

A system quad fetch will then occur, in order to bring in the remainder of the 16-byte line. This quad fetch can execute concurrently with any later read/write hit cycles driven by the CPU.

To avoid any coherency or ownership problems in the period of time after the CPU write has occurred and before the system fetch has begun, the write into the cache memory 72 RAM array section 100 (the write miss) and the following system quad fetch must execute as one indivisible operation.

The architecture of the controller 70 guarantees that the local write cycle and the following system quad fetch will in effect be a locked operation (although SLOCK# is not asserted). To ensure this, the write miss will not be completed until the controller 70 owns the system bus (SHLDA deasserted). Write misses when the system bus is granted will be stalled by delay of RDYO# until SHLDA is deasserted. If the controller 70 owns the system bus when the write miss occurs, the controller 70 will perform the system quad fetch to completion, by not granting SHLDA until the fetch and any subsequent possible write-back cycles are completed.

Similar to read misses, write misses are of two types: write line misses and write tag misses:

1. Write line misses:

A cacheable write line miss will latch the write data into the cache data array, and assert RDYO# in three

- 94 -

clocks to the CPU to terminate the cycle on the host side. In addition, a system quad fetch will be initiated in order to obtain the remaining data of the cache line. Incoming data from the system quad fetch will be merged
5 with any dirty data in the line, including the just-written miss data from the CPU. No write-backs will occur from write line misses, as no tags are replaced. The details are as follows:

Signal MALE is asserted in T2 to inhibit writes
10 directly into the cache data array. The data is instead written first into write register 120, and then into the data array. A "mask" bit associated with each doubleword of the line inside the cache memory 72 is set. This
"mask" bit is used to prevent the data written from the
15 486 CPU and any other dirty data from being overwritten by the corresponding data which will later be obtained from the pending quad fetch.

RDY0# is asserted to the 486 CPU to terminate the local CPU cycle. This termination happens during the
20 third clock of the cycle (one wait state).

The host port of the cache memory 72 can serve any read/write hits with no wait states following any write miss. Figure 46 shows an example of a write miss followed directly by a read hit, and the subsequent processing on
25 both busses that occurs.

A write miss cycle detected by the controller 70 results in the miss address being latched into the cache memory 72 miss address register 110. The controller 70 then generates a quad fetch to main memory, at the
30 location pointed to by miss address register 110. This quad fetch is "wrapped around" the demand miss word on a 16-byte address boundary. The order will either follow 486-style address order or sequential address order.

Data returned from the quad fetch is loaded into the
35 cache memory 72 memory update register set 116. Similar to system quad fetches due to read line misses, the DW# signal is asserted from the controller 70 to the cache

- 95 -

memory 72's to qualify incoming quad fetch data, so that the write miss data and any other dirty data is not overwritten.

The controller 70 asserts SBLAST# during the fourth 5 transfer of the quad fetch, to terminal the cycle. At the completion of the system quad fetch, the controller 70 then writes the contents of memory update register set 116 into the cache memory 72 data array by assertion of the QWR (Quad Write) signal. The update entry in the data 10 array is pointed to by the miss address register 110.

Figure 47 shows a write line miss being terminated with one wait state, and the resulting system quad fetch that occurs.

2. Write tag misses:

15 Write tag misses are similar to write line misses. However, a tag will be replaced during write tag miss processing, so write tag misses may be followed by write-back cycles. The quad fetch will be followed by one or two write-back cycles of the data for the replaced tag, if 20 either line contained any "valid" or "dirty" data. Write-back cycles due to write tag misses are otherwise identical to write-backs from read tag misses.

First, MALE is asserted by the controller 70 in T2 to prevent the write data from directly overwriting data in 25 the cache array, since this data may be dirty and must then be written back to memory. Instead, the write data is internally latched into the cache memory 72 write register 120.

Next, a line to be replaced is selected. The 30 controller 70 will select one out of the two selected replace entries if any of the entries are marked "invalid" of 2-way set associative mode has been chosen. The LRU bits and an LRU policy will be used to select if all entries of the replace lines are marked "valid". The 35 selected line is latched into the cache memory 72 write back register set 118.

Now the data in write register 120 is written into

- 96 -

the cache memory 72 data array, since any potentially dirty data has been moved to write back register set 118. As described before in the write line miss case, there is a mask bit for each doubleword to prevent the CPU write data from being overwritten by the incoming system quad fetch data.

The controller 70 then terminates the write miss cycle through RDYO# assertion. As in the write line miss case, this will occur in three clocks.

10 The write miss address is latched into miss address register 110, and a system quad fetch occurs at the address pointed to by the miss address register 110, with the fetched data loading memory update register set 116. As before, SBLAST# terminates the system quad fetch.

15 Assertion of QWR by the controller 70 writes the contents of memory update register set 116 into the cache data array.

A subsequent read/write operation from the 486 with a different miss line address will be 'frozen' through delay of RDYO# or BRDYO# until the replacement cycle (if any) of this write miss is completed.

Finally, one or two write-back cycle(s) will occur if any valid and dirty data was replaced. These write-backs will occur as follows:

25 Write-back cycles due to write tag misses are functionally identical to the write-backs generated from read tag misses which were previously described. System logic can and should not differentiate between these two cases of write-backs as they serve the same purpose and function identically.

Figure 48 shows a write tag miss. The line retrieved from the resulting system quad fetch replaces dirty data and a write-back cycle to main memory is generated.

SYSTEM BUS STATES

35 All system cycles will begin in the ST1 state, when the controller 70 asserts SADS# (low) and drives the cycle

- 97 -

definition and address signals. Bus cycle definition and encodings will emulate that of the i486 CPU, with the exception that SBLAST# will be driven to a valid level in ST1, along with the other definition signals. All 5 controller 70 signals (except one) are driven to valid levels in the ST1 state.

The ST2 state will always follow the ST1 state. The controller 70 will remain in the ST2 state until either SRDYI# or SBRDYI# is asserted. The DW# output of the 10 controller 70 will be asserted only in ST2 states. DW# will not be valid in ST1 states.

If the memory system chooses a non-burst transfer by returning SRDYI# for a multiple-cycle transfer, the controller 70 will re-enter the ST1 state and drive 15 another ADS# signal to the system.

Once either SRDYI# or SBRDYI# is asserted to complete the first cycle in a multiple-cycle transfer, the same signal must again be asserted to complete the remaining cycles, or correct operation is not guaranteed. For 20 example, if the system asserts SBRDYI# to complete the first cycle of a system quad fetch, SBRDYI# must be asserted to complete the remaining cycles of the transfer. This requirement is the same for both multiple cycle read and write transfers. Assertion of SBLAST# by the 25 controller 70 indicates that the next assertion of the ready inputs will terminate this cycle.

In some cases, System Hold Acknowledge (SHLDA) latency can be longer than the corresponding hold acknowledge (HLDA) latency for the i486 CPU. SHOLD will 30 be acknowledged by asserting SHLDA only if no further cycles are required to complete the previous transaction. For example, if a write miss triggered a system quad fetch further followed by two write-back cycles, an SHOLD request occurring during the write miss would not be 35 acknowledged until the write-back cycles had completed on the bus.

- 98 -

OTHER BUS OPERATIONS

This section will detail how the controller 70 will respond to operations other than 486 CPU bus cycles. Reset, flush, assertion of SHOLD, and snoop operations 5 will be covered.

RESET

RESET clears the tag arrays (valid and dirty bits) of the controller 70, RESET should not always be tied directly to the system, as some systems may desire to 10 reset the CPU separatably from the memory array. Unlike write-through caches, the controller 70 must be reset with the memory array in these cases. The controller 70 will be able to respond to SHOLD during RESET, as the 486 recognizes HOLD during RESET.

15 SHOLD/SHLDA OPERATION

The controller 70 grants and receives the system side bus through the use of SHOLD (System Hold Request) and SHLDA (System Hold Acknowledge). SHOLD may be asserted at any time to the controller 70. If the system bus is idle, 20 SHLDA will acknowledge the request by asserted high. However, if a system bus cycle is in process, there will be a latency until the request is acknowledge. If SHOLD is still active at the next system bus cycle boundary (SBLAST# low and either SRDY# or SBRDY# asserted low), 25 SHLDA will then be asserted high in acknowledgement. As in the 486 CPU, short SHOLD requests (those that appear during a bus cycle bus disappear before the first bus cycle boundary) are ignored.

In the same clock that SHLDA is asserted, the 30 controller 70 will float all system side control, address, and data lines to allow another bus master access to the system bus.

Operation of the control signals for the optional address transceivers is shown in Figure 49. These 35 transceivers are turned around when the controller 70

- 99 -

either grants or receives the bus. This turning is transparent when the controller 70 grants the bus, as it occurs during the idle state that SHLDA is asserted, so systems need not wait to drive addresses when they sample 5 SHLDA high.

Due to the one clock latency of turning the transceivers, an extra Ti state will occur when the controller 70 receives the bus. However, the system needs no redesign to account for this. Since this one-clock 10 latency is guaranteed, the SHOLD request may be relinquished one clock early in order to avoid any system performance degradation.

FLUSH OPERATIONS

In 80486 systems, flush operations can be invoked in 15 two ways, one software and one hardware.

The first method of generating flushes is by software execution of the 486 INVD and WBINVD instructions correspondingly. The 486 CPU correspondingly generates special cycles during these instructions, in addition to 20 flushing its internal cache. The INVD produces the Flush cycle, while the WBINVD produces the Write-back cycle. The encodings for these cycles were previously shown.

Because the controller 70 utilizes a write-back architecture, its response to both Flush and Write-back 25 cycles is identical. For both cycles, the controller 70 first copies all lines marked "dirty" back to memory. The duration required to execute the write-back cycles depends on the number of "dirty" lines to be copied back to main memory. The latency for controller 70 to complete this 30 cycle will be much longer than that of the 486 CPU, due to these write-back cycles. Because of this latency, the controller 70 will recognize SHOLD requests during flush operations. SHLDA will be granted at the completion of the write-back cycle(s) corresponding to the tag entry 35 currently being flushed. Next, the controller 70 will clear (flush) all directory valid bits, LRU bits, and

- 100 -

dirty bits.

Note that system memory controllers will never explicitly see the Flush and Write-Back cycle encodings. These cycles are intended for external caches, and the controller 70 does not pass them on to the system.

The hardware method of generating a flush is through the activation of the FLUSH# input signal. The controller 70 FLUSH# pin should be connected directly to that of the 486 CPU.

Assertion of the FLUSH# pin is treated as if a Write-back cycles has been received, with the controller 70 writing back dirty data to memory. Since the HOLD/HLDA protocol is used before flush operations are commenced, the 486 CPU will not produce another bus cycle before this write-back is completed. Once the flush operation is completed, HOLD will be deasserted to the 486 CPU and normal operations will resume.

The latency of flush operations will be quite long, since all lines which are partially dirty must be written back to main memory. However, failure of some system devices may occur if the controller 70 requires use of the system bus for such lengthy intervals. Because of this, the controller 70 will acknowledge SHOLD with SHLDA during flush operations. If SHOLD is asserted to the controller 70 during a flush operation, SHLDA will be returned at the end of one or two write-back cycles for the tag entry which is currently being flushed. When SHOLD is released, the controller 70 will resume writing back dirty lines to main memory until completed.

30 WRITE-BACK CYCLES DUE TO FLUSH OPERATIONS

As discussed, write-back cycles to memory will be performed when either the FLUSH# pin is asserted or either of the FLUSH or Write-back special cycles occur. Flush operations are the second way that write-back cycles can be generated. As previously shown, normal read/write tag misses may also produce write-back cycles.

- 101 -

For each line of the data cache that contains valid dirty data, the entire line will be written to memory in a quad doubleword Operation. Since there are two lines associated with each tag entry, either zero, one or two quad writes will occur for each tag. If two writes occur for a tag (both lines contain valid and dirty data), the controller 70 will assert MWB at the end of the first write in order to inform the cache memory 72 that the second line for the tag must be written to memory as well.

Figure 50 details the beginning of a hardware flush operation due to assertion of the FLUSH# pin. The controller 70 responds by acquiring the bus from the 486 CPU and begins the process of writing back lines which contain valid dirty data to memory. Flush Write-back cycles are generally identical to write-backs due to read/write tag misses, described before. DW# will be asserted for each doubleword to indicate to the system if the doubleword is dirty or not. System logic should use DW# as described in the Read Miss section, in both write enable logic and to hasten assertion of SRDYI#/SBRDYI# for non-dirty dummy writes.

The controller 70 must be granted the local processor bus, in order to perform the flush operations. To acquire the 486 CPU local bus, the controller 70 will use the HOLD/HLDA protocol for either hardware or software flush operations. If the FLUSH# pin is asserted (hardware flush), the controller 70 will immediately assert HOLD to the CPU. When HLDA is returned by the CPU, the controller 70 will begin the flush operation, by using the local CPU bus to drive out flush addresses (addresses of lines which contain dirty data) and asserting CALE to latch these addresses within the cache memory 72.

Although the concept is the same, software-generated flush operations signalled by the Flush and Write-back cycles will begin slightly differently. As these operations are signalled by 486 CPU bus cycles, the controller 70 will assert HOLD to the CPU, followed by

- 102 -

RDY0#. As the CPU recognizes HOLD on bus cycle boundaries, HLDA will be driven at the end of the 486 CPU Flush or Write-back cycle. Having obtained the bus, the controller 70 begins driving out flush addresses and CALE 5 as previously described.

SNOOP OPERATIONS

The system memory controller must accommodate the controller 70 when snoop operations occur on the system bus due to another bus master, whether by DMA or another 10 CPU. The controller 70 supports snoops through the SEADS# pin, which is functionally equivalent to the 486 CPU EADS# pin. In addition, the controller 70 has three new pins not present on the 486 CPU to allow snoop operations to correctly occur: SNPBUSY, SMEMWR#, and SMEMDIS.

15 As the controller 70 has no SAHOLD (System Address Hold) pin, the controller 70 system bus must be tri-stated through the SHOLD/SHLDA protocol before the snoop can occur. However, this is not a limitation for single-CPU processing, as the SHOLD/SHLDA method is the simplest and 20 preferred method of allowing DMA to occur.

As SEADS# is asserted, the controller 70 will sample SMEMWR# to determine if a snoop read or write is occurring. In either case, the SNPBUSY output toggles high in acknowledgement.

25 SNOOP WRITES

The combination of SEADS# low and SMEMWR# high informs the controller 70 that a snoop write is in progress on the system bus. Snoop writes will trigger two responses by the controller 70.

30 First, the controller 70 will obtain the local CPU address bus through assertion of either AHOLD and/or BOFF#. This will allow the controller 70 to drive an invalidation cycle to the i486 CPU at the snoop write address. EADS# will be asserted to the CPU to trigger 35 this invalidate internally in the CPU.

- 103 -

Second, SNPBUSY will be asserted high to acknowledge the snoop operation. The controller 70 will check its internal array to determine if the snoop write is a hit or a miss. If the address is present and valid in the cache
5 (a hit), the controller 70 will write the data for the snoop write directly into the cache memory 72 RAM array section 100.

This will occur if the snoop write is a byte, word (2-byte) or doubleword write, and regardless of whether
10 the data is marked as dirty or clean. This write appears as a local bus operation. No change in the tag array dirty or valid bits will occur. There is no output of the tag lookup to inform the system if a hit or miss has occurred for a given snoop write.

15 On a snoop write, the system may process the write normally. There are two requirements that must be met in order for the controller 70 to correctly process the snoop write:

1. Since a snoop write hit will be updated into the
20 cache data array, the system must supply valid data and byte enables on the controller 70's system wide for all snoop writes. This data and byte enables should be driven valid along with the system address SA<27:2>, SMEMWR#, and SEADS# pins.
- 25 2. As previously stated, SNPBUSY toggles high in response to a snoop operation. SNPBUSY will remain high for some number of clocks, while the snoop operation is continuing. The falling edge of SNPBUSY indicates to the system that the controller 70 has
30 completed its internal lookup and is prepared to terminate the snoop write operation.

The system must return SRDYI# or SBRDYI# after SNPBUSY has fallen, in order to terminate the snoop write operation. Since the external system cannot determine if
35 a hit or miss has occurred, SRDYI# or SBRDYI# must be returned for all snoop write operations. If the snoop write has not completed on the system bus when the falling

- 104 -

edge of SNPBUSY occurs, the system is free to withhold SRDYI# or SBRDY# from the controller 70 and add wait states to the operation. SRDYI# and SBRDYI# are not sampled by the controller 70 during the snoop write operation before SNPBUSY has fallen. As a result, the system is free to hold SRDYI# or SBRDYI# low while SNPBUSY is high, in anticipation of terminating the cycle as quickly as possible upon the falling edge of SNPBUSY.

SNOOP READS

- 10 The controller 70 detects snoop reads through the combination of SEADS# low and SMEMWR# low. When this happens, SNPBUSY toggles and stays high for a varying number of clocks while the tag lookup occurs. While this tag lookup is in progress, the system memory controller should delay the cycle until one or two cases occurs:
- 15 1. The requested data is not contained in the cache data array (a snoop read miss), and the system memory must supply the data. This result is indicated by SNPBUSY going low while SMEMDIS has remained low. On the clock that SNPBUSY is sampled low, the system memory controller may determine that it needs to supply the requested data. From the clock in which SEADS# is asserted to the controller 70, latency of performing the snoop read operation will be two clocks for the snoop miss case. A snoop read miss is shown in Figure 51.
- 20 2. The requested data is contained in the cache RAM array and is 'dirty' (a snoop read hit), resulting in the need to supply the requested data to the system bus. SMEMDIS (System MEMORY DISable) is the resulting output which will be asserted in this case. A snoop read hit will be indicated to the system by the activation of SMEMDIS (high), while SNPBUSY is still high.
- 25 3. The clock after SMEMDIS is activated, SNPBUSY will be de-asserted and the cache memory 72 enabled by assertion

- 105 -

of SP OE# to drive the requested data onto the system bus. The snoop address will be passed to the cache memory 72 via the local CPU bus.

To acquire the 486 CPU local bus from the CPU, the controller 70 asserts AHOLD followed by assertion of the BOFF# output. BOFF# will be released two clocks after the snoop operation has terminated. Any local 486 CPU bus cycle that is aborted in this manner will be restarted by the CPU, another BOFF# is released. The controller 70 generated the snoop address and asserts both CALE and MALE in T1, latching this address inside the cache memory 72 miss address register 110.

The cache memory 72 will continue to be driven until the system asserts SRDYI# to the controller 70. SRDYI# is sampled beginning in the clock in which SMEMDIS is activated. SRDYI# being asserted incorrectly in this clock will cause the controller 70 snoop cycle to terminate, even though no snoop data has yet been driven onto the bus. A snoop read hit operation is shown in Figure 52.

The maximum frequency allowable for snoop writes and reads is different. Due to the basic architecture of the 486 CPU and isolation of the system bus from the local CPU bus, snoops can never be performed on consecutive clocks. Snoops to the controller 70 may be performed as fast as every third clock if the snoops are consecutive snoop writes, or a snoop write followed a snoop read. A snoop write takes three clocks to be processed internally by the controller 70. Three clocks after a snoop write, another snoop of either type may be performed. In addition to performing an internal lookup and possible invalidation, the controller 70 will pass the snoop address and EADS# on to the 486 CPU. From the clock that SEADS# is asserted, there is a latency of three clocks to pass the address and EADS# to the 486 CPU for either the hit or miss case.

Because of the latency of accessing the data in the cache memory 72's snoop reads are necessarily slower than

- 106 -

the snoop writes. After a snoop read occurs, the system should not assert another snoop read until SRDYI# is asserted, terminating the present snoop. It is allowable to begin another snoop read in the same clock that SRDYI# is returned.

PARITY

The controller 70 architecture supports the parity function of the i486 CPU. The bytes in the cache memory 72 data array have x9 structure, in order to support parity.

HOST PORT PARITY

Since the i486 CPU generates even parity on write cycles, the cache memory 72 store the parity bit(s) that is driven as the CPU write occurs. When a CPU read hit occurs, the parity bit(s) will be returned to the CPU, along with the requested data.

CACHE MEMORY 72 SYSTEM PORT PARITY

As on the host port 113, the cache memory 72 have no parity generation logic on the system port. For cycles in which data is read into the cache data array from the system port side, it is the responsibility of the system to supply proper parity bits. This applies for both normal read cycles and snoop write cycles.

For normal read operations on the system port, it is sufficient for the system to provide storage for the parity bits, and return them on read cycles. This is allowable, since even parity will be driven from the cache memory 72 write cycles to the system port, whether induced by write-backs or CPU writes which are passed on to the system.

Since data is written into the burst-RAM data arrays during write hits, the system should provide even parity on the data pins during snoop write operations.

- 107 -

CONTROLLER 70 HOST POST PARITY

In addition, the controller 70 supports parity for the accesses that occur to its control registers. Unlike the cache memory 72, however, the cache memory 72 dynamically generates its parity bit during read cycles from these registers. This generation is necessary since the control register contents may change without any explicit write cycles from the CPU. As an example, the CPU may set the Expansion Register INV bit in order to flush the cache. At completion of the flush, the controller 70 will clear this bit in order to resume normal operation.

Referring back to Figure 32, the sequencing states of concurrent bus control unit 206 of cache controller 70 are next considered. The sequencing states will be described with reference to Figures 53-56.

Referring first to Figure 53, the sequencing of the concurrent bus control unit 206 begins when the ADS# signal is received by cache controller 70 (step 300). When the ADS# signal is received, state 302 is entered. At state 302, ADDR<14:4> is latched into the hit address register 109 of cache memory 72. The tag address is compared with the received address, and the line valid bits are checked. It is noted that the address array information ADD ARY indicates the address information stored within controller 70. Furthermore, the attribute array ATT ARY stores line valid, tag valid and dirty bit information, and the IRU ARY array stores bank select information.

If the received address information is a hit read operation as determined by steps 304 and 306, the concurrent bus control unit 206 enters state 308. At state 308, the data from memory array 100 corresponding to the hit address is latched into memory read hold register set 114 one clock later. In addition, controller 70 triggers a burst read operation by asserting BRDY# to the CPU HPOEA#, HPOEB# to the burst-RAM to trigger a read

- 108 -

operation from the memory array. Following state 308, the sequencing of concurrent bus control unit 206 returns to its start state, and is available for additional hit operations. LRU array inside cache controller will be
5 updated at the end of state 308.

If the operation is determined to be a write hit operation at steps 304 and 306, the sequencing enters state 310. At state 310, the data at host port 113 is latched into memory write register 120, and a write
10 operation is triggered from memory write register 120 to the RAM array 100. In addition, the controller asserts the signals HW0# and HW1# and the line dirty bits are updated. LRU array data will be updated at state 310. The write in BRAM cache memory 72 is self-timed. Write
15 actions inside cache memory 72 will not start until the next clock after RDY is returned to CPU.

If the operation is determined to be a miss operation by step 304, the sequencing goes to step 312 as shown in Figure 54. If the operation is a read miss operation as
20 determined by steps 312 and 314, state 316 is entered. At state 316, signal ADDR<14:4> is latched into miss address register 110. In addition the corresponding data from memory array 100 is latched into memory write back register set 118. The controller 70 asserts signals MALE,
25 HPOEA#, HPOEB#, WBSTB, and updates the LRU array. State 318 is thereafter entered. The cache controller 70 asserts signal SADS# and updates the ADDR array, and the signal bypass is asserted by controller 70 and data is thereby caused to pass from the system port 112 to the
30 host port 113. Following state 318, the data is latched into memory update register set 116 when signal SBRDY# is returned, during bus state 320. In addition, the associated valid bits are set. The controller 70 begins the burst read from the system and resets all tag and line
35 valid bits. In addition, controller 70 reads and stores all dirty bits of the replaced line. State 322 is next entered and the subsequent data of the read miss burst

- 109 -

sequencing is latched into memory update register set 116, when signal SBRDY# is returned, and the associated valid bits are set. In addition, controller 70 starts the burst read from the system. At state 324, a write operation
5 from memory update register set 116 to memory array 100 is triggered. The controller 70 asserts signals QWR, MWB if the stored dirty bit indicates that the next line of the current replaced block contains dirty data, (if TN4D = 1), and the tag valid and line valid bits are updated for the
10 fetched line.

At step 326, it is determined whether the current line contains dirty data. If the current line contains dirty data, state 328 is entered wherein the contents of memory write back register set 118 are dumped to the
15 system port 112. The controller 70 asserts signal SADS#, SPOE# (first replacement). Step 328 is bypassed if the current replaced line contains no dirty data. At step 330 the next line is checked to determine whether it contains dirty data. If the next line contains dirty data, state
20 332 is entered and the data is latched from memory array 100 into memory write back register 118 after the miss address pointer is bumped by prior MWB assertion through controller 70. In addition, the data in memory write back register set 118 is provided to system port 112.
25 Controller 70 furthermore asserts signals SADS#, SPOE#, MWB (second replacement).

It is apparent from Figure 54 that if either the current line or the next line do not contain dirty data, one or both of states 328 and 332 will be bypassed.
30 Following these states, the sequencing is ready to receive further ADS# signals. However, other hit operations can start processing concurrently after state 324.

If a write tag miss is determined at step 314, state 340 is entered as shown in Figure 55. At state 340,
35 ADDR<14:4> is latched into miss address register 110. In addition, the corresponding data in RAM array 100 is latched into memory write back register set 118, and the

- 110 -

data at the host port 113 is latched into memory write register 120. Controller 70 asserts signals MALE, HPWEA#, HPWEB#, and WBSTB. The LRU array is furthermore updated. The corresponding mask bit is set during the next clock
5 after state 340 and the host port data is written into the burst-RAM array section 100. State 342 is thereafter entered wherein the data is started to pass from the system port 112 to the host port 113. A write operation from memory write register 120 to RAM array 100 is
10 furthermore triggered. Furthermore, the corresponding mask bit is set. Controller 70 asserts signals SADS#, RDY#, and the address array is updated. At state 344, data is latched into memory update register set 116 when signal SBRDY# is returned, and the associated valid bits
15 are set. Controller 70 starts a burst read from the system, and the tag valid and line valid bits are reset. Dirty bits of the replaced line are read and stored. Next, at state 346, the succeeding data of the burst quad fetch read is latched into memory update register set 116
20 when signal SBRDY# is returned, and the associated valid bits are set. Controller 70 furthermore starts a burst read from the system. Finally, at state 348, a write operation from memory update register set 116 to RAM array 100 is triggered, and controller 70 asserts signals QWR
25 and MWB if the stored dirty bit indicates that the next line of the current replacing block contains dirty data (if TN4D = 1). Furthermore, the tag valid, line valid, and dirty bits are updated for the fetched line.

If the current line is determined to contain dirty
30 data at step 350, state 352 is entered. The data in memory write back register set 118 is provided to system port 112, and the controller asserts signals SADS# and SPOE# (first replacement). Similarly, if the next line is determined to contain dirty data at step 354, state 356 is
35 entered wherein the data from RAM array 100 is latched into memory write back register set 118. Furthermore, the data in memory write back register set 118 is provided to

- 111 -

the system port 112. Controller 70 asserts SADS#, SPOE# (second replacement). Subsequent operations may thereafter occur since the sequencing returns to step 300. However, other hit operations can start processing 5 concurrently after state 348.

If a line miss operation is determined at step 312, steps 360 and 362 determine whether either a read line miss or a write line miss has occurred. If a read line miss has occurred, the sequencing enters state 364. At 10 state 364, signal ADDR<14:4> is latched into memory address register 120. Controller 70 asserts signals SADS#, MALE, HPOEA#, HPOEB# and BYPASS. Furthermore, the LRU array is updated. State 366 is thereafter entered wherein the data is latched into memory update register 15 set 116 when signal SBRDY# is returned. Furthermore the associated valid bits are set. Controller 70 accordingly triggers a burst read operation. At state 368, a write operation from memory update register set 116 to RAM array 100 is triggered, and cache controller 70 asserts signal 20 QWR. The line valid bits are furthermore updated.

On the other hand, if a write line miss occurs, state 370 is entered. Signal ADDR<14:4> is latched into miss address register 110. In addition, the host port data is latched into memory write register 120 and the associated 25 mask bit is set. Controller 70 furthermore asserts signals SADS#, MALE, and HPWEA#, HPWEB#. The LRU array is furthermore updated. State 372 is thereafter entered wherein a write operation from memory write register 120 to RAM array 100 is triggered, mask bit is set for 30 advanced writes, and the data is latched into memory update register set 116 when signal SBRDY# is returned. In addition, the associated valid bits are set. Controller 70 asserts signal RDY# and a burst read is triggered. Finally, at state 374, a write operation from 35 memory update register set 116 to RAM array 100 is triggered, and controller 70 asserts signal QWR. The line valid and dirty bits are furthermore updated.

- 112 -

Bus controller 200 of cache controller 70 as shown in Figure 32 is next considered with reference to Figure 57. Bus controller 200 controls CPU 60 in accordance with the state machines 400 through 410 as depicted in Figure 57.

5 The initial state 400 is entered when no operations between the CPU 60 and cache memory 72 occur.

For a 386 based system, if signal ADS# is asserted, the state machine of bus controller 200 enters state 402. Following the sequencing as described above within the
10 concurrent bus control unit 206, the signal RDY# is asserted and the state machine of bus controller 200 returns to its initial state 400.

For a 486 based system, for a write operation, when signal ADS# is asserted, the state machine of bus
15 controller 200 enters state 404. Upon processing of the 404 state by concurrent bus control unit 206, signal RDY# is asserted and the state machine returns to state 400.

In the case of a burst read operation, after state 404 has been processed, signal BRDY# is asserted and the
20 state machine of bus controller 200 enters state 406. The appropriate processing by concurrent bus control unit 206 thereby is initiated and bus states 408 and 410 are sequentially entered. Following execution of machine state 410, the state machine returns to state 400.

25 Referring next to Figure 57, the state machines of controller 202 (Figure 32) are next described. The initial state is state 420. An operation in write through mode causes the state machine to go to state 430. This state change is initiated by signal TSREQ and the
30 deassertion of signal TSBURST. Upon completion of state 430, signals SRDY#/SBRDY# are asserted. The state machine accordingly resets to state 420.

For signal TSREQ and TSBURST assertions state 422 is entered. The state machine is returned to state 420 upon
35 assertion of SRDY# or SBRDY# for noncacheable or scalar reads. For a burst read operation, signal SBRDY# is asserted when the state machine is at state 422.

- 113 -

Subsequent assertion of SBRDY# causes the states to change from state 422 to state 428. Upon completion of the burst read, the state machine returns to its initial state 420.

State 432 is entered if external bus master wishes to own the system bus. System bus ownership is granted upon entering into state 432 and signal SHOLDA is asserted. Once signal SHOLD is deasserted, state machines state 202 returned to initial state 420.

From the above description, it should be noted that the burst RAM cache memory chip in accordance with the invention may also support cache systems with line sizes larger than four doublewords. During a replacement cycle, activation of an external signal by cache control logic causes an internal counter to access neighboring lines of data so that lines of eight or 16 doublewords can be saved for burst write-back operations transparent to system memory 38. This activity does not prevent the read miss data from being fetched first, eliminating what would be a large replacement time penalty.

It is also noted that port 113 interfaces with the Intel 386 or 486 processor in a manner similar to conventional SRAMs. If data parity is not checked as in most 386 systems, HP<8> and SP<8> can be left disconnected.

By decoupling the host and system data buses 112 and 113, processor cache accesses and system memory accesses can proceed simultaneously. Decoupling on a write-back cache allows write misses to be handled with zero wait states since read and write cache hits can proceed in parallel with the write miss data fetching. Moreover, when combined with the procedure implemented by each burst RAM memory chip for saving write-back data during read or write misses, the dual-port architecture allows write-back line replacement cycles to be completely hidden from the processor.

When used in sets of four, the burst RAM memory chips allow internal data transfers of four doublewords

- 114 -

(128 bits) on reads or writes, achieving a true 128-bit quadword transfer in only one clock. When doing memory burst operations, the internal structure of the burst RAM memory chip behaves similar to a static column RAM in that
5 the input buffer and word line delays occur only on the first access to the first subarray. Subsequent accesses in the burst are supplied by successive subarrays which are pre-addressed. This structure allows full speed burst accesses without the use of 10 nsec SRAMs.

10 The architecture of burst RAM cache memory 72 eliminates the historical disadvantage of write-back caches: having to do a replace cycle before doing the data fetch. With burst RAM cache memory 72, the replace data is saved within a set of latches of the burst RAM memory
15 chips (memory write back registers 118A-118D) as soon as a miss is detected and thus allows the fetch to begin immediately. Only after the fetch is complete does a replacement cycle begin, and during that time, the host processor 60 can resume accessing the cache RAM array 100
20 through the host port 113. As an example, in a system with a two wait state memory and non-burst fetches, the use of burst RAM cache memory 72 allows the 386 to be operating at full speed after a read miss in just 5 clocks. Using standard SRAMs, the same system would take
25 48 clocks (32 clocks for two lines on a tag miss plus 16 clocks for the new data).

Decoupled buses also avoids the necessity to redesign the system memory subsystem every time the processor module is upgraded. In fact, the processor data bus and
30 the system data bus can run at different speeds. Since the processor will work out of its cache subsystem more than 95% of the time, the system memory need not be redesigned to run at the same speed as the processor. This is especially true if the system uses a burst memory
35 controller so that the cache memories can support burst operation to and from system memory. The main memory system might then operate at a reasonably easy to design

- 115 -

speed of 20 MHz, while the processor module (processor, cache controller, and cache memories) operates at 25, 33, or even 40 MHz. Upgrading the processor from an 80386 to an 80486 would involve only minor design changes to main 5 memory.

Also noted above, both of the 9-bit (eight bits plus parity) data ports 112 and 113 of each burst RAM memory chip support "demand word first" wrapped around quadword operations as well as scalar reads and writes. The system 10 port 112 supports burst operations to and from system memory 38 if the system memory controller also supports burst.

The dual-port architecture of the burst RAM cache memory chip permits processor and main memory accesses to 15 occur in parallel, while hiding write-back cycles from the processor, contributing to a substantial performance increase over alternative implementations. The burst RAM cache memory chip fully supports an 80386 using a write-back cache, including support for non-cacheable accesses 20 and multiple replacement cycles for lines longer than 16 bytes.

The cache memory 72 also supports quadword data fetch and, though not common in 386 systems, burst operation with main memory. If the cache and system memory 25 controllers support burst operation, cache memory 72 will also support burst reads and writes to system memory 38. In an 80386 system which is designed for future upgrade to an 80486, the main memory subsystem may use Intel's strongly suggested 64-bit, bank interleaved main memory 30 organization. This memory organization is accessed using the non-sequential burst order used by the 486. Cache memory 72 directly supports that burst order to and from main memory, as well as the sequential burst order used with the smaller, less expensive, and standardized 32-bit 35 sequential memory organization.

If the system memory uses the design-intensive 64-bit, bank interleaved architecture recommended by Intel

- 116 -

instead of the standard, 32-bit sequential architecture, then the system port 112 uses the 486 burst sequence to system memory 38 instead of sequential order. Burst RAM cache memory 72 supports both sequential and 486 burst
5 ordering. In either case, the fetch data from main memory is brought in "demand word first" such that the first doubleword passes directly to the 386 through each corresponding bypass path 119 of cache memory 72. The 386 can then resume execution while the remainder of the burst
10 data is brought in and stored in the RAM array 100.

Numerous modifications and variations will become apparent to those skilled in the art once the above disclosure is fully appreciated. It is to be understood that the above detailed description of the preferred
15 embodiment is intended to be merely illustrative of the spirit and scope of the invention and should not be taken in a limiting sense. The scope of the claimed invention is better defined with reference to the following claims.

- 117 -

CLAIMS:

I Claim:

1. A memory cache apparatus comprising:
a random access memory;
5 a host port;
a system port;
an input latch coupled to said host port for selectively writing data to said memory; and
an output register coupled to said system port
10 for receiving data from said memory and selectively furnishing said data to said system port.
2. An apparatus as in Claim 1, further comprising an input register connected to said system port for furnishing data to said memory.
- 15 3. An apparatus as in Claim 2, wherein said input latch is a memory write register, said input register is an update register, and said output register is a write back register for furnishing data to said system port.
- 20 4. A memory cache apparatus comprising:
a random access memory;
a host port;
a system port;
an input register coupled to said host port and to said random access memory for selectively writing
25 input data to said random access memory; and
an output register coupled to said system port for receiving output data from said random access memory and selectively furnishing said output data to said system port;
30 wherein said input data is written into said random access memory from said input register at the same time when said output data is furnished to said system port from said output register.

- 118 -

5. The memory cache apparatus as recited in Claim 4 further comprising:

5 a miss address register coupled to said random access memory for storing an addressing signal corresponding to said output data; and
a hit address register coupled to said miss address register for storing an addressing signal corresponding to said input data.

6. The memory cache apparatus as recited in Claim 1
10 further comprising a second output register coupled to said random access memory and to said host port for furnishing data from said memory to said host port.

7. The memory cache apparatus as recited in Claim 6 wherein said second output register is a read hold
15 register.

8. The memory cache apparatus as recited in Claim 6 wherein said input register can store a plurality of words of data.

9. The memory cache apparatus as recited in Claim 8
20 further comprising means for masking the writing of selected words of data into said random access memory.

10. The memory cache apparatus as recited in Claim 1 further comprising a bypass path coupled between said host port and said system port for directly allowing the
25 passage of data between said host port and said system port.

11. The memory cache apparatus as recited in Claim 5 further comprising a counter coupled to said miss address register.

30 12. The memory cache apparatus as recited in Claim

- 119 -

11 wherein said counter increments from an address
associated with a first line of data to a subsequent line
of data.

13. The memory cache apparatus as recited in Claim 2
5 further comprising means for validating data stored within
said input register to selectively prevent the writing of
data of said data stored within said input register into
said random access memory.

14. The memory cache apparatus as recited in Claim 1
10 wherein said random access memory includes a plurality of
parity bits.

15. The memory cache apparatus as recited in Claim 1
wherein said RAM is organized in a plurality of lines
wherein each of said lines comprises a plurality of word
15 storage locations, and wherein each of said word storage
locations is selectively writable.

16. The memory cache apparatus as recited in Claim 1
wherein said random access memory is single ported.

17. The memory cache apparatus as recited in Claim
20 16 wherein said random access memory has a wider bandwidth
than said host port and said system port.

18. The memory cache apparatus as recited in Claim 1
wherein said random access memory performs a read-modify-
write operation.

25 19. A method for operating a memory cache apparatus,
said memory cache apparatus including a random access
memory, a host port, a system port, an input register
coupled to said host port, and an output register coupled
to said system port, said method comprising the steps of:

- 120 -

latching input data into said input register
from said host port;

comparing a received address to a plurality of
cache addresses;

5 loading replaced data from said random access
memory into said output register if said received
address does not match one of said plurality of cache
addresses;

10 loading said input data into said random access
memory; and

providing said replaced data to said system
port.

20. The method for operating a cache memory
apparatus as recited in Claim 19 further comprising the
15 step of loading subsequent input data into said input
register at the same time that said replaced data is
provided from said output register to said system port.

21. The method for operating a cache memory
apparatus as recited in Claim 19 wherein data from a
20 plurality of data locations of said random access memory
are provided to said output register during a single clock
cycle.

22. A computer system comprising:
a host microprocessor having a host address bus
25 and a host data bus;
a system memory having a system address bus and
a system data bus;
a dual port cache memory having a system port
connected to said system data bus and a host port
30 connected to said host data bus; and
a cache controller connected to said cache
memory.

23. The computer system as recited in Claim 22

- 121 -

wherein said cache controller provides a first address on said host address bus at the same time said cache controller provides a second address on said system address bus, said first address corresponding to a
5 different memory location than said second address.

24. The computer system as recited in Claim 22 wherein data on said host data bus is asynchronous to data on said system data bus.

25. The computer system as recited in Claim 22
10 wherein said cache controller comprises:
a first control sequencer for controlling addressing and data signals on said host address bus and on said host data bus; and
a second control sequencer for controlling
15 addressing and data signals on said system address bus and on said system data bus.

26. The computer system as recited in Claim 22 further comprising means for disabling said dual port cache memory during a local bus access cycle.

20 27. The computer system as recited in Claim 22 further comprising a peripheral device coupled to said system memory.

28. The computer system as recited in Claim 27 wherein said peripheral device provides data to said
25 system data bus, and wherein a hit address memory location within said dual port cache memory is loaded with said data from said peripheral device if the hit address of said dual port cache memory corresponds with an address of said data from said peripheral device.

30 29. The computer system as recited in Claim 22 wherein said host microprocessor operates at a first

- 122 -

frequency, and wherein said system memory operates at a second frequency that is different from said first frequency.

30. The method for operating a cache memory
5 apparatus as recited in Claim 19 further comprising the step of coupling said host port to said system port when a read miss cycle occurs.

31. The method for operating a memory cache
apparatus as recited in Claim 19 wherein said cache memory
10 apparatus further comprises an update register for providing data from said system port to said random access memory, said method comprising the further step of loading update data from said system port into said update register.

15 32. The method for operating a memory cache apparatus as recited in Claim 19 wherein said step of latching input data into said input register from said host port occurs in a first clock cycle and wherein said step of loading said input data into said random access
20 memory occurs on a second clock cycle, said second clock cycle immediately following said first clock cycle.

1/70

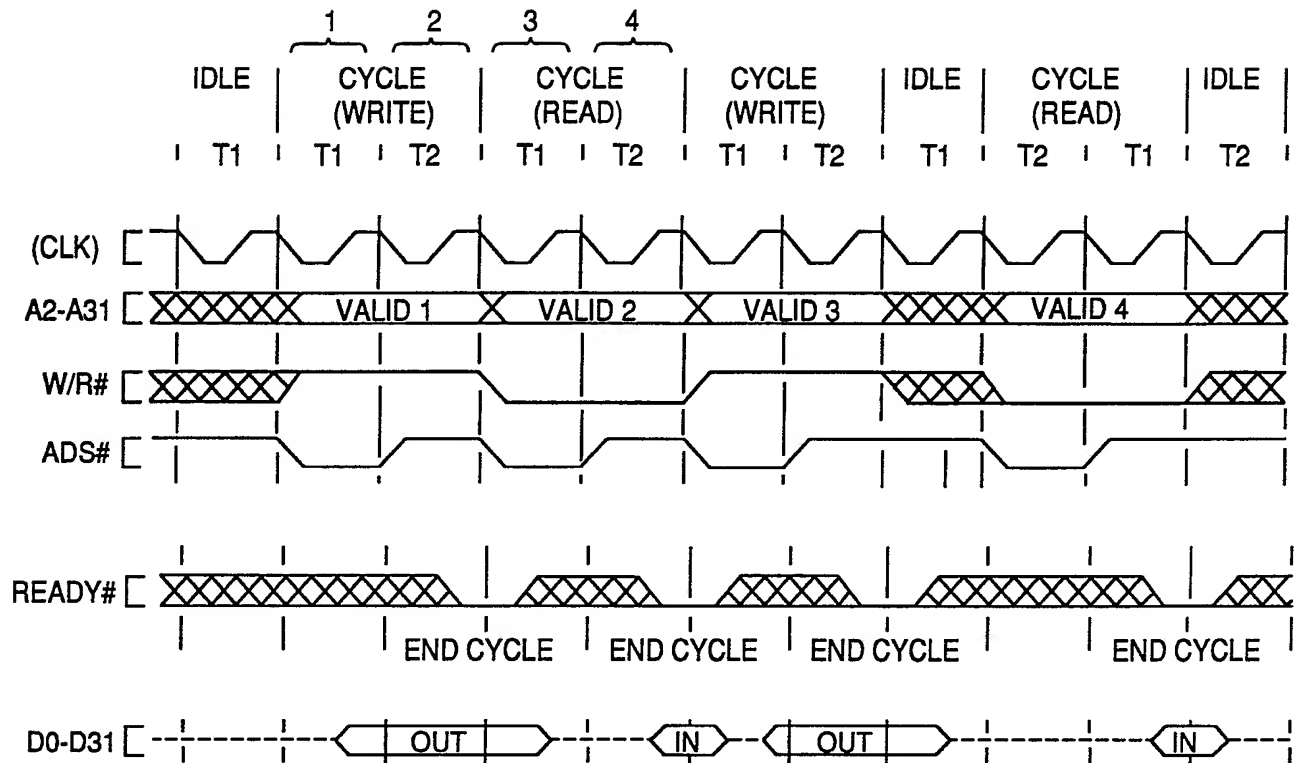


FIG. 1
(PRIOR ART)

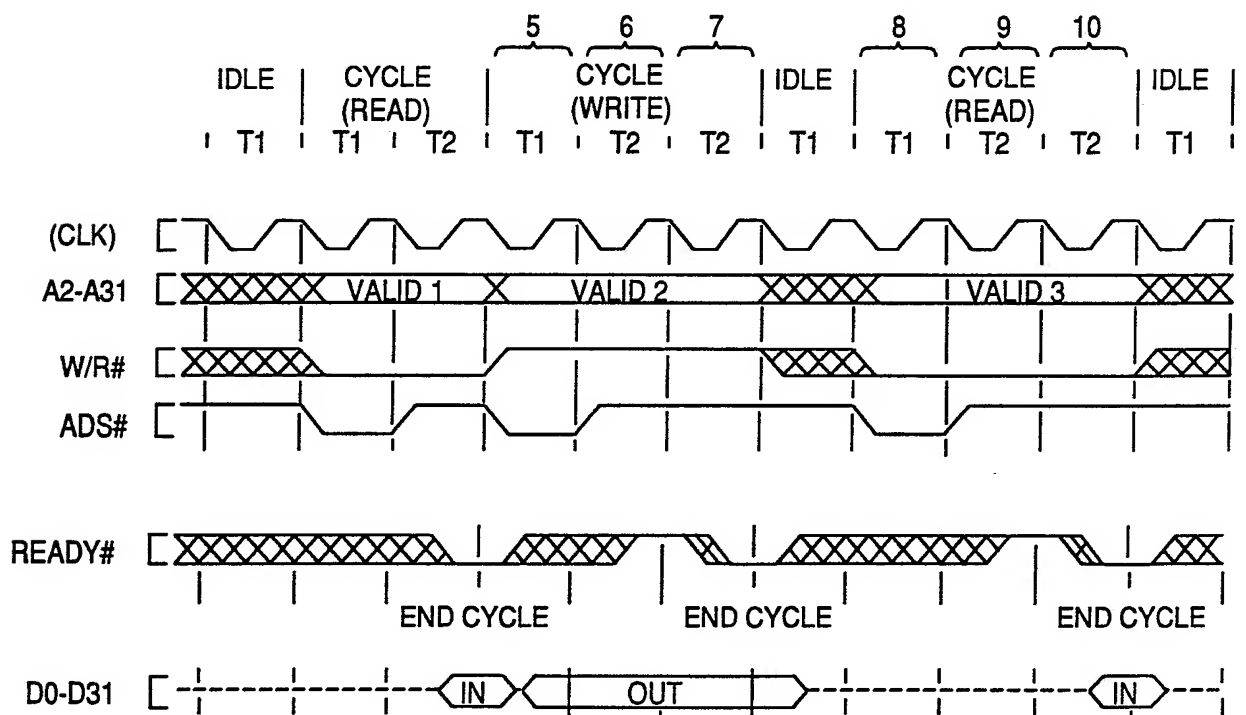


FIG. 2
(PRIOR ART)

2/70




FIG. 3
(PRIOR ART)





FIG. 5
(PRIOR ART)

3/70

FIG. 4
(PRIOR ART)

4/70




FIG. 6

5/70



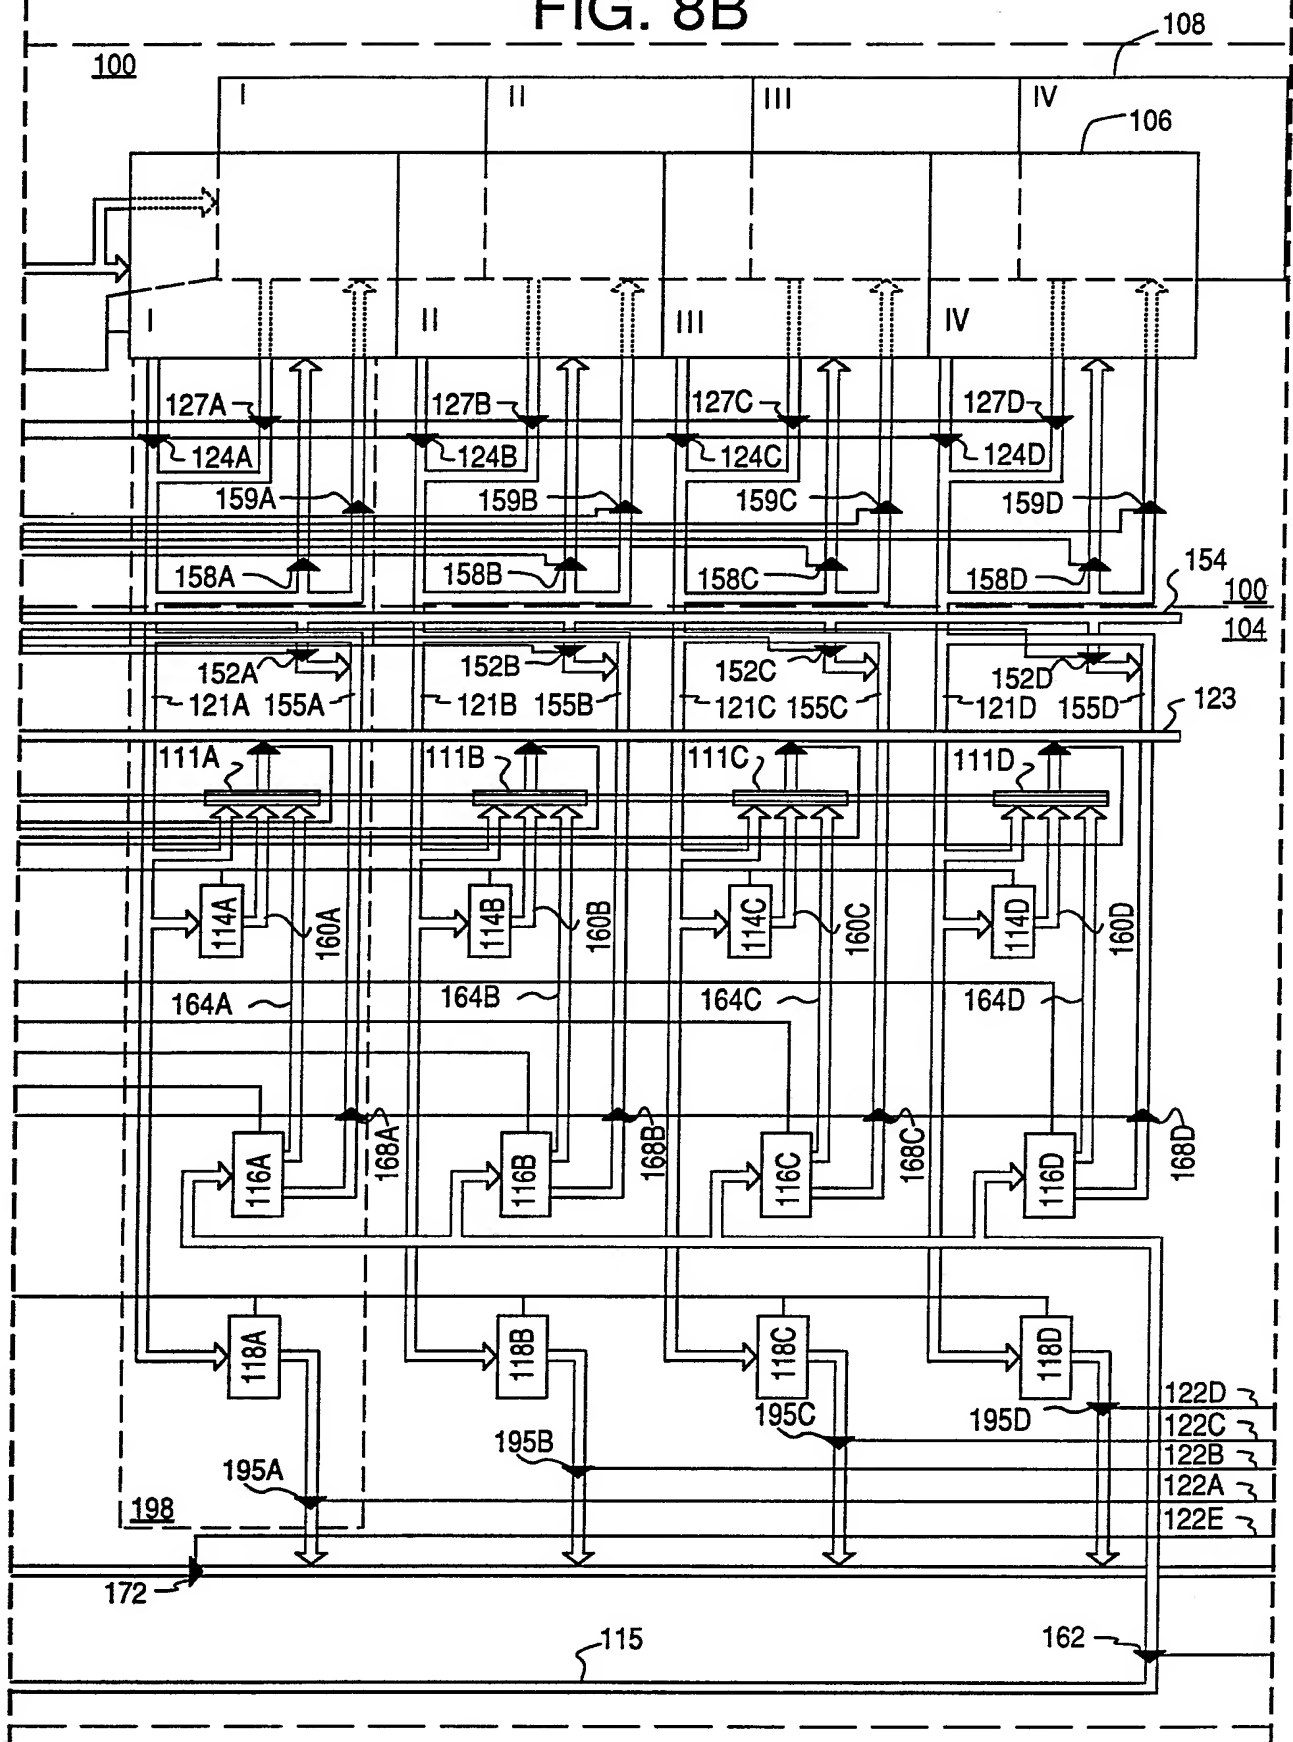


FIG. 7

6/70

FIG. 8A



7/70
FIG. 8B

8/70





FIG. 8C

9/70



10/70




FIG. 9A'

11/70





FIG. 9A''

12/70

FIG. 10



13/70




FIG. 10A'




FIG. 10A"

15/70




FIG. 11

16/70


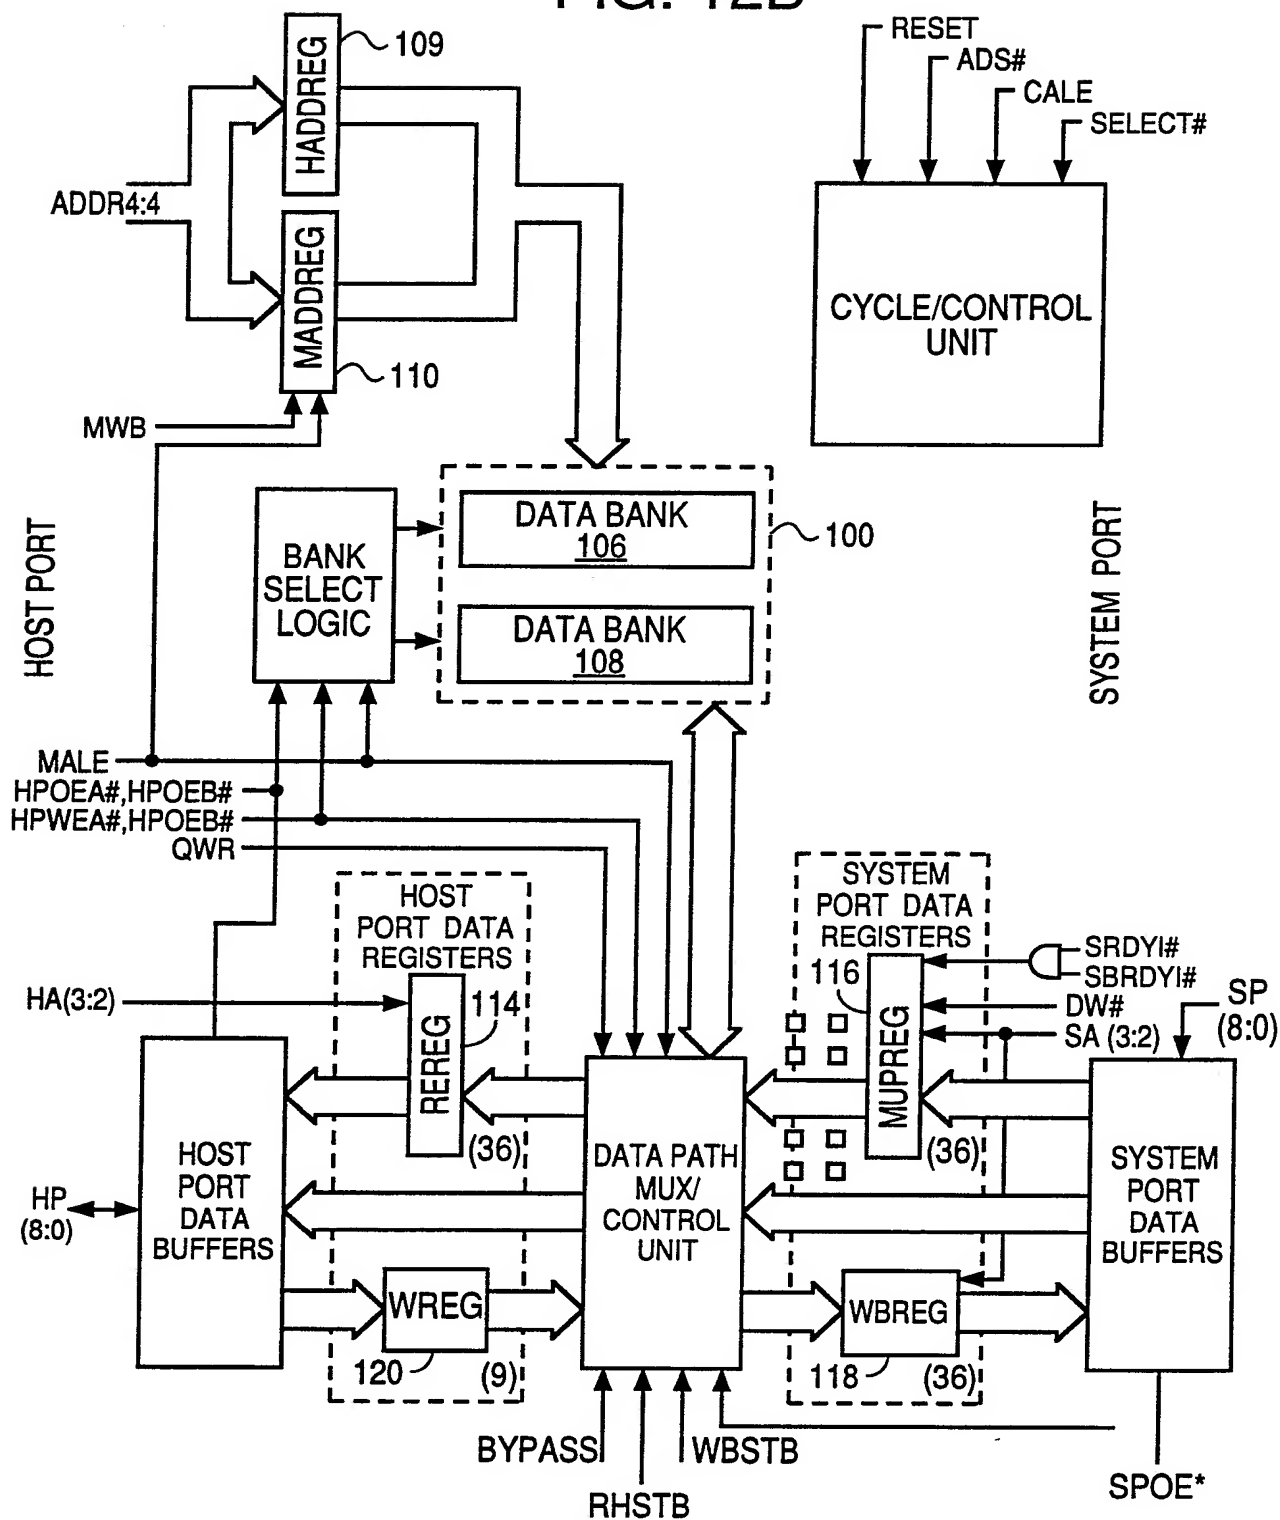


FIG. 12A


17/70

FIG. 12B



18/70

FIG. 13



19/70

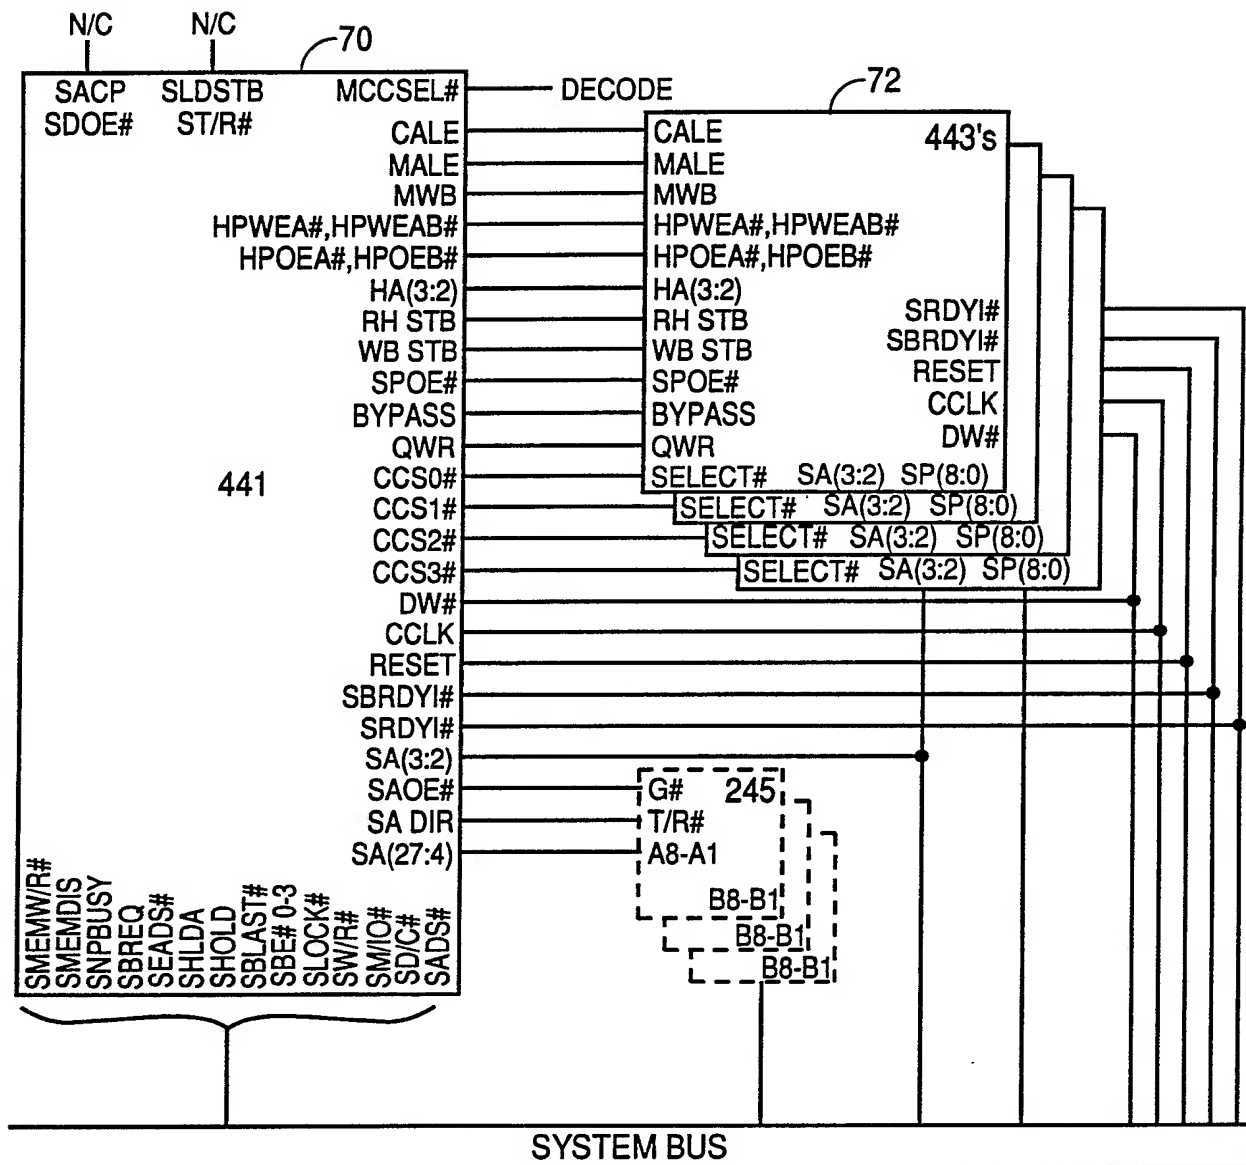


FIG. 14

20/70


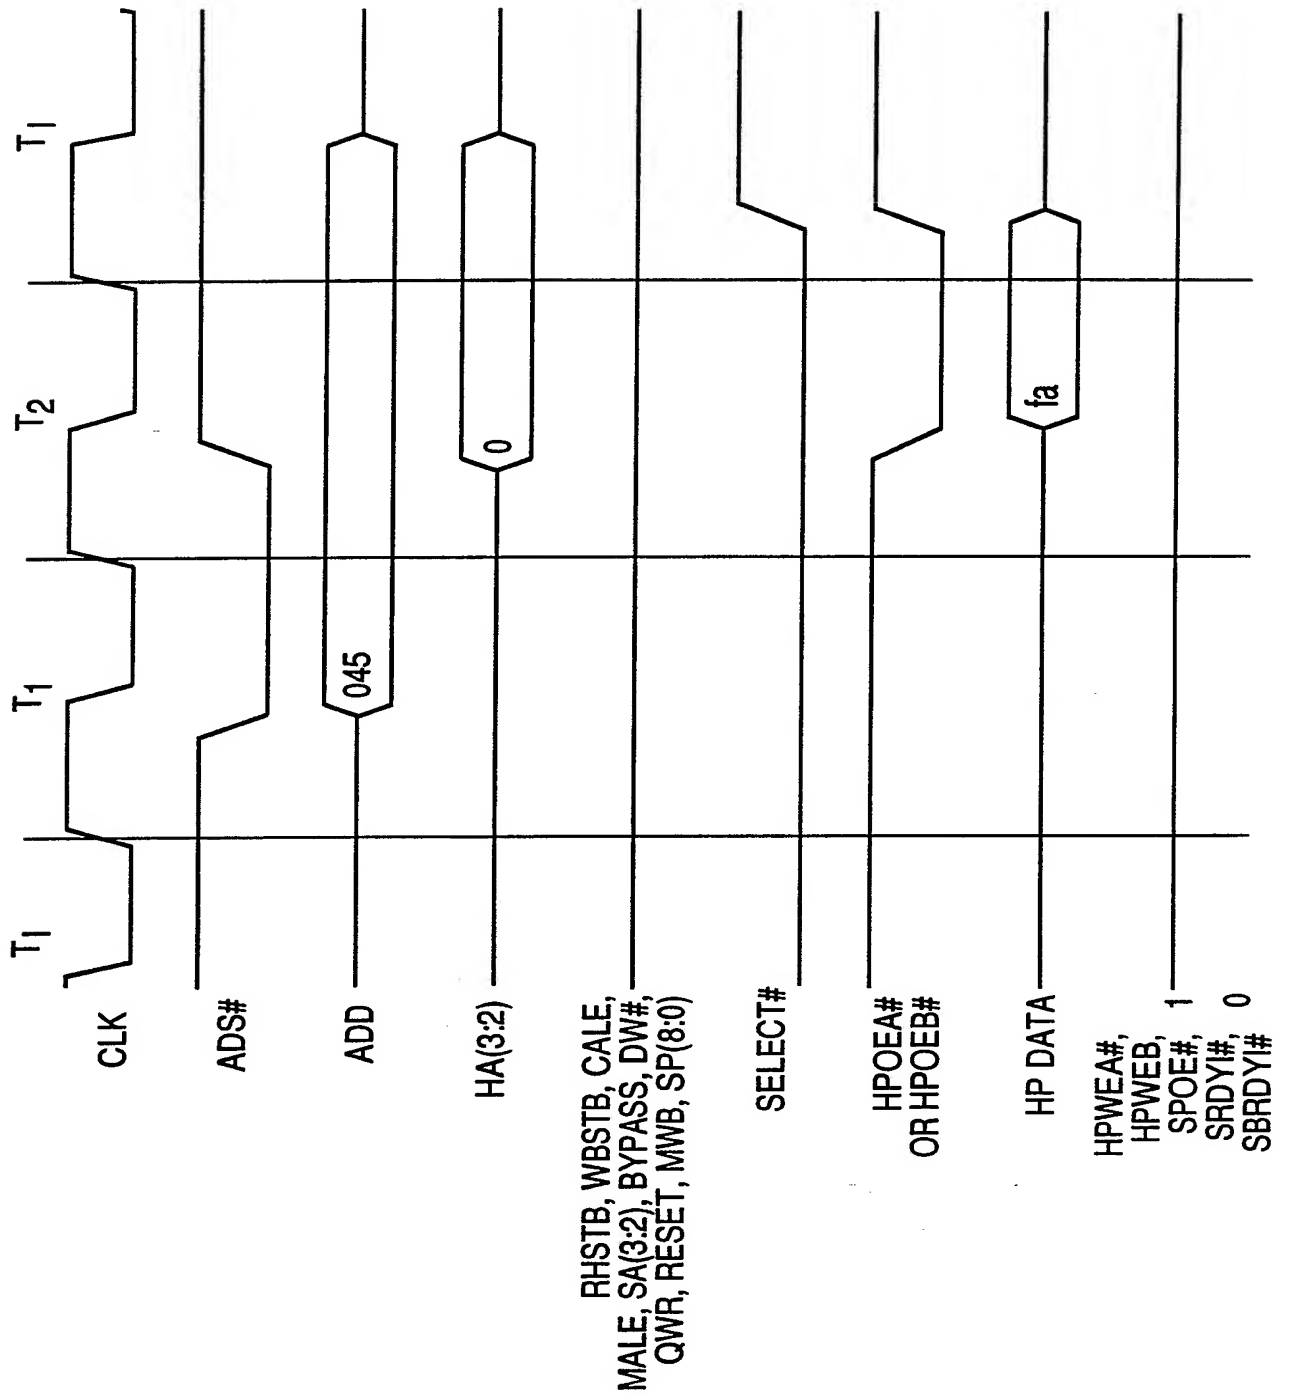


FIG. 15

21/70

FIG. 16



22/70

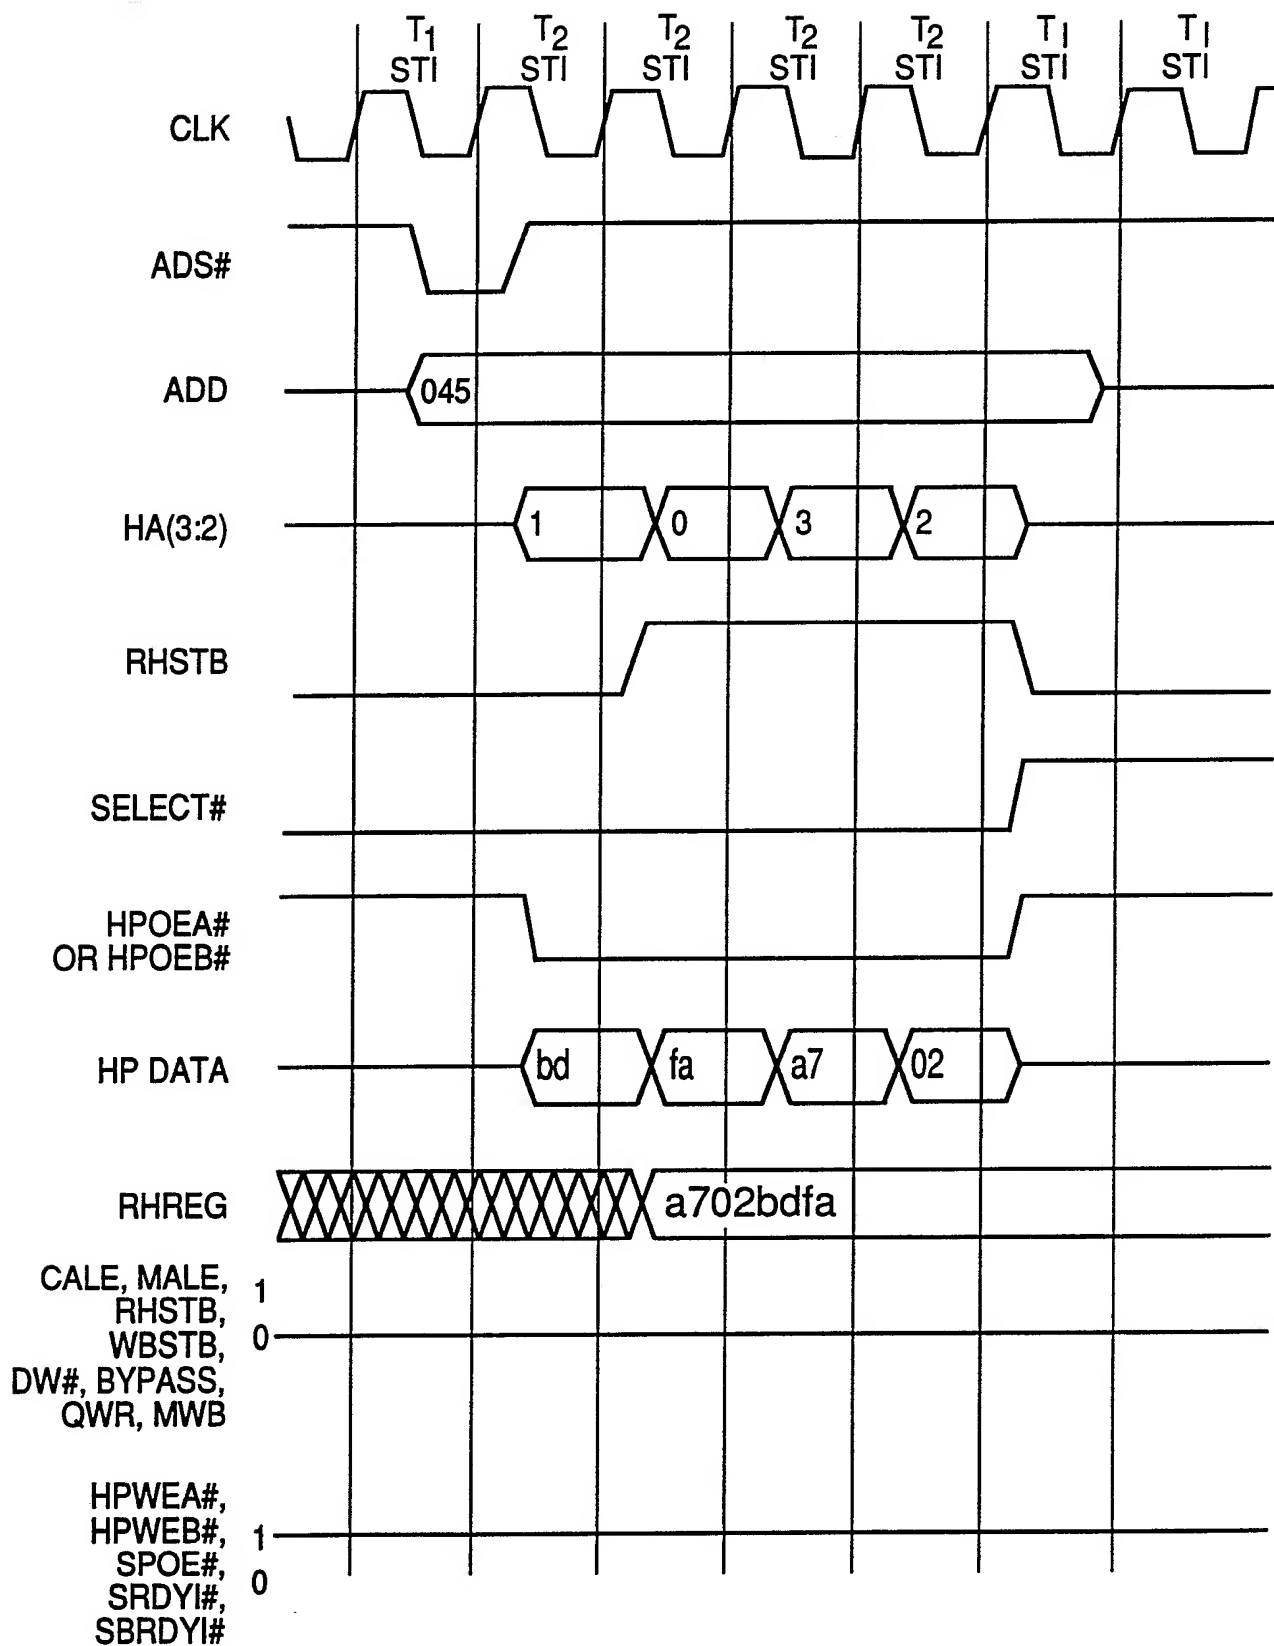


FIG. 17

23/70




FIG. 18

24/70

FIG. 19




FIG. 20


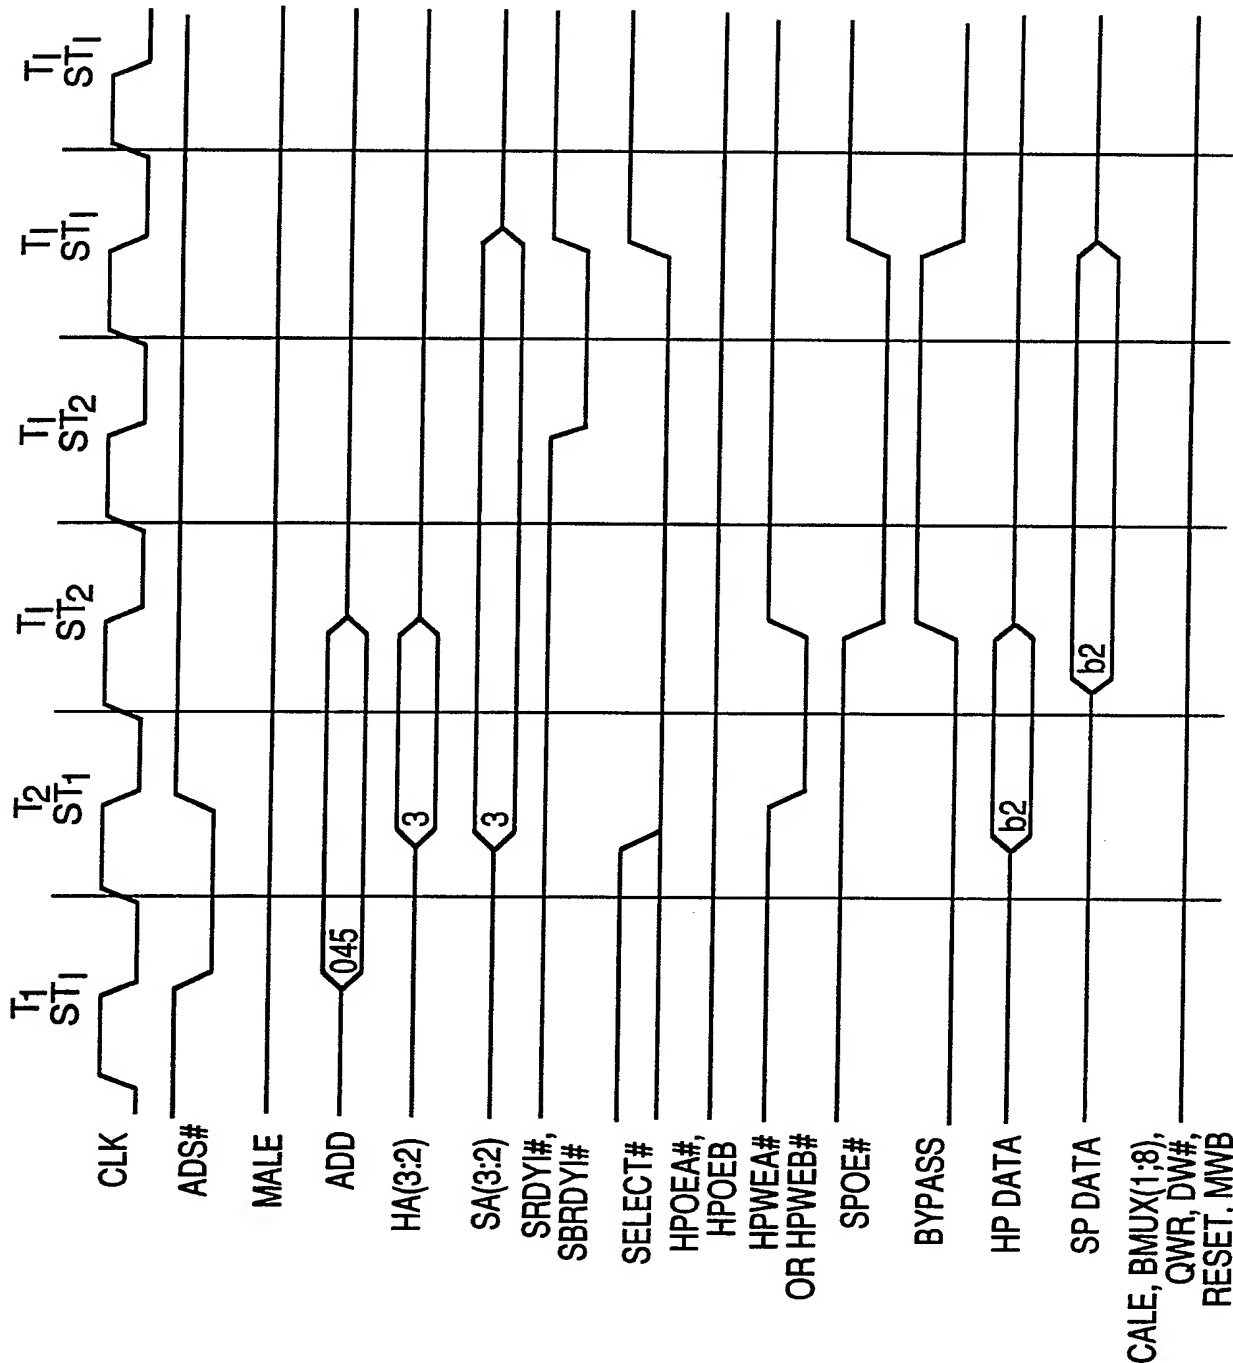
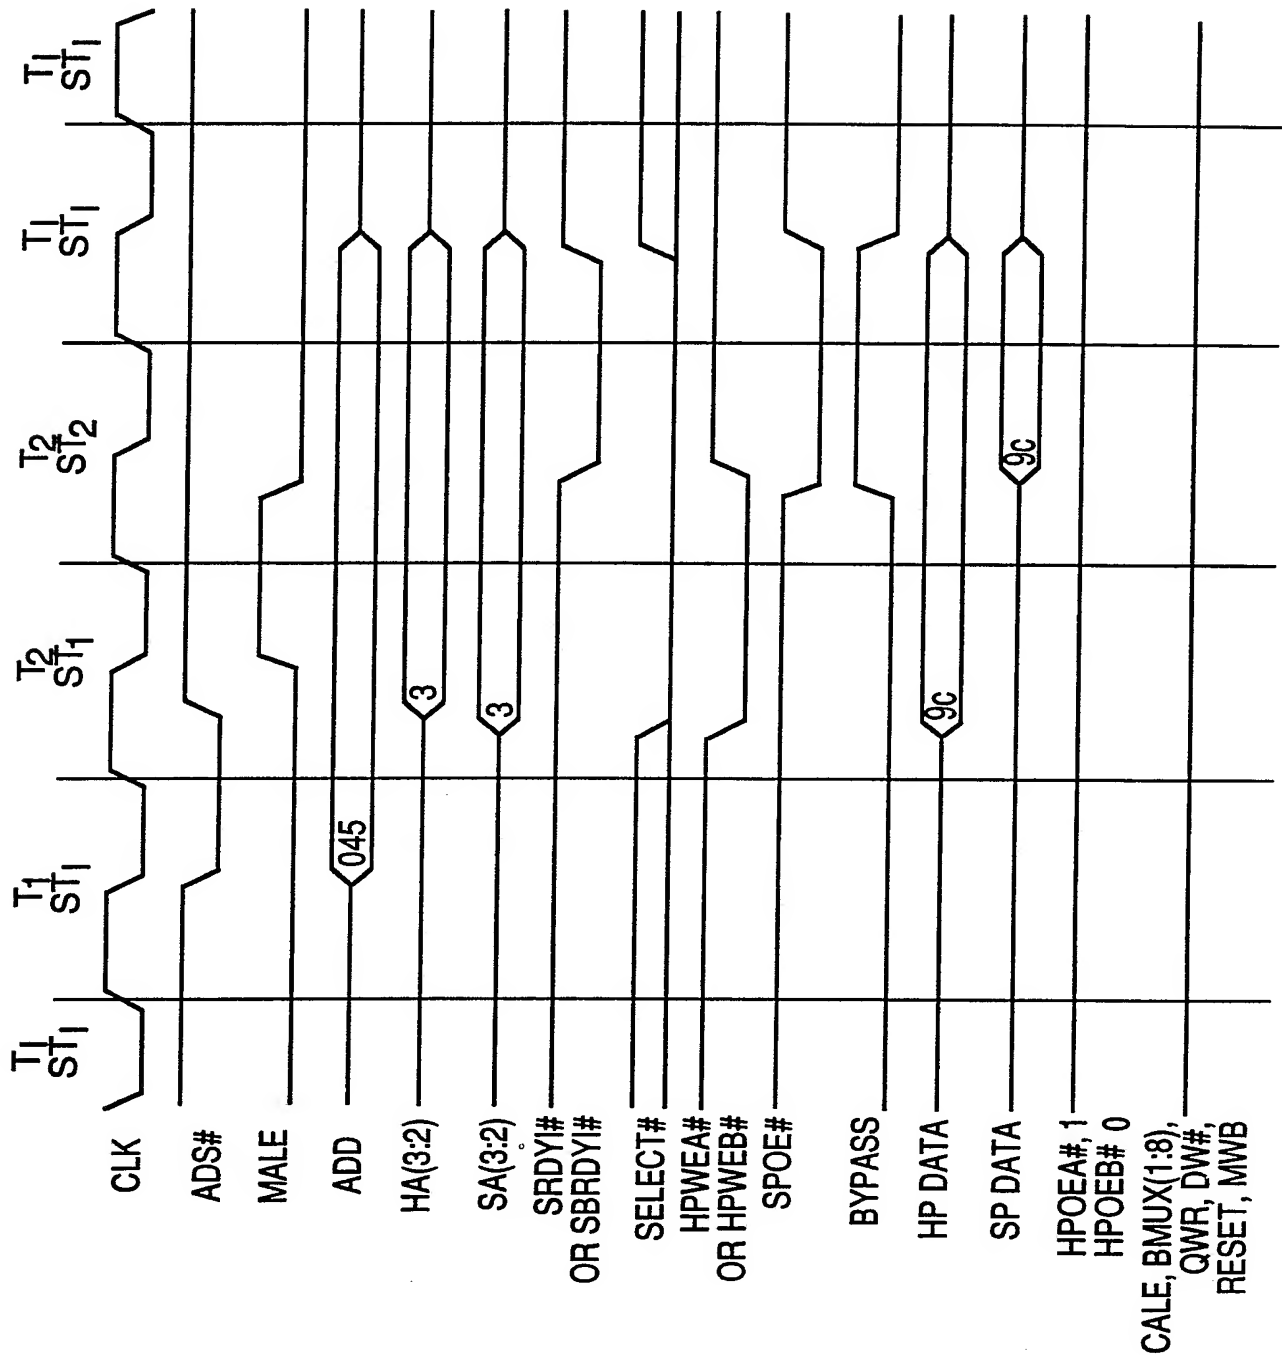


FIG. 21



27/70

FIG. 22



28/70

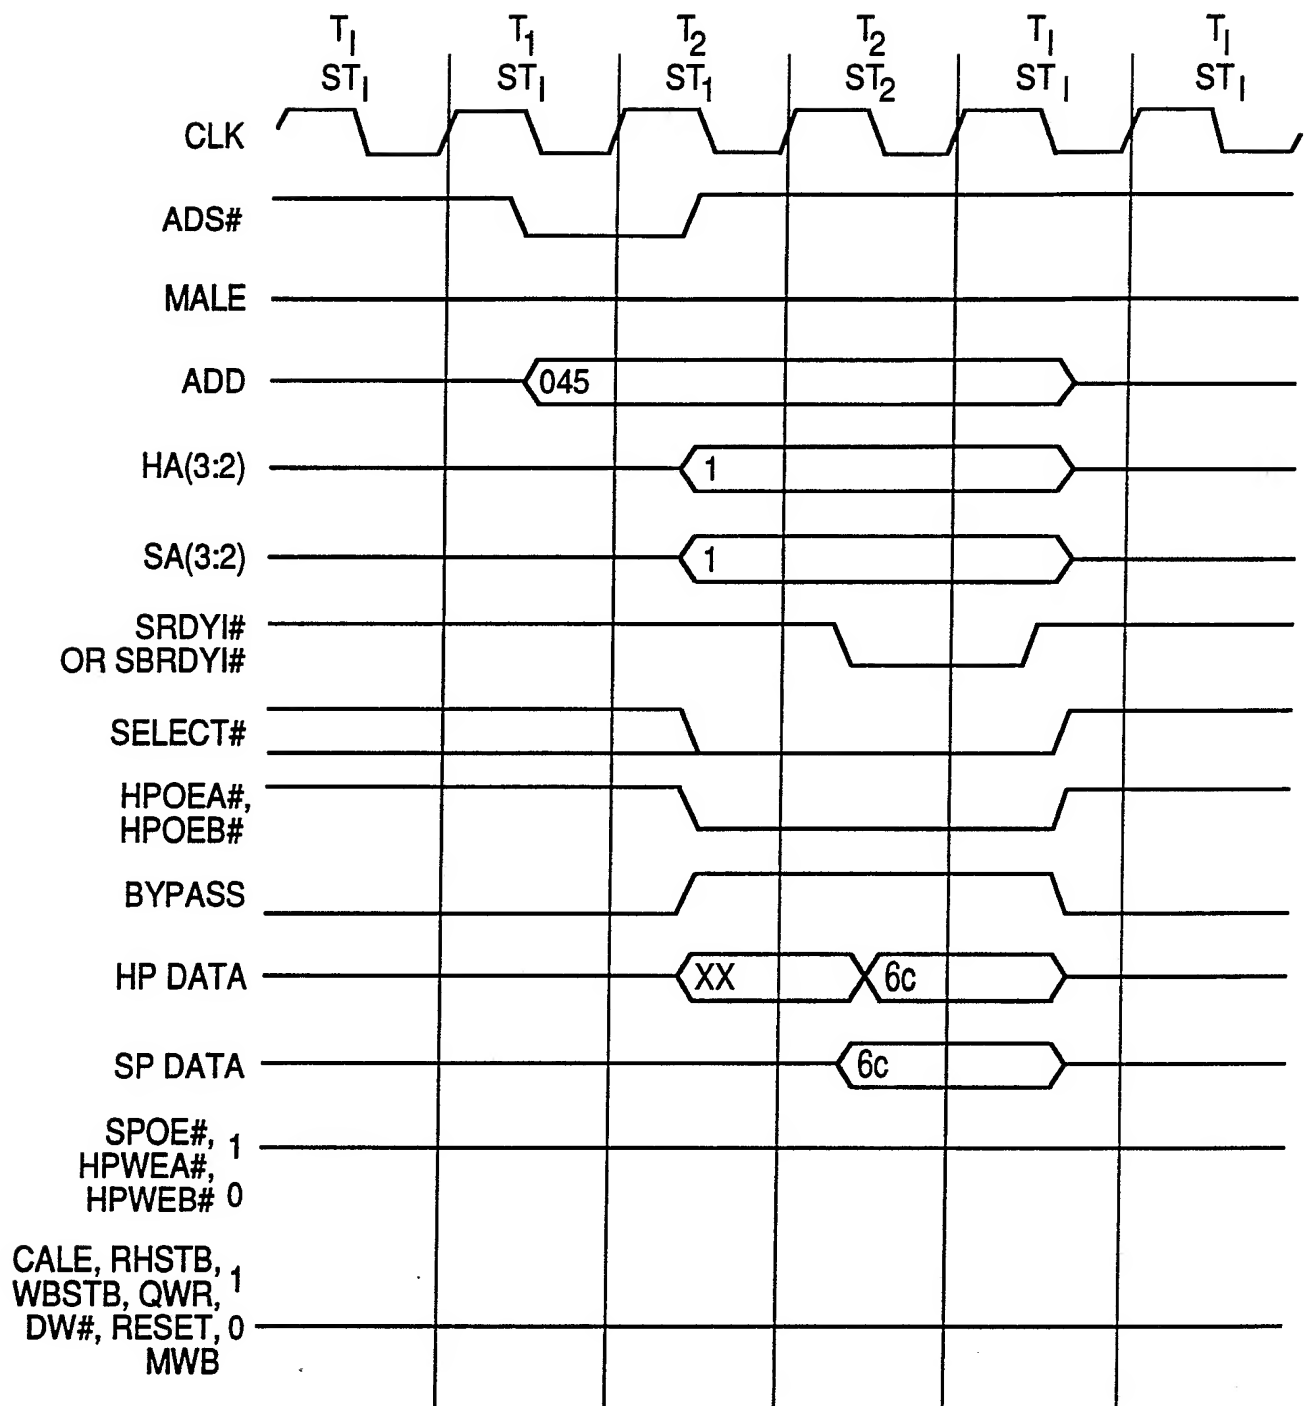


FIG. 23

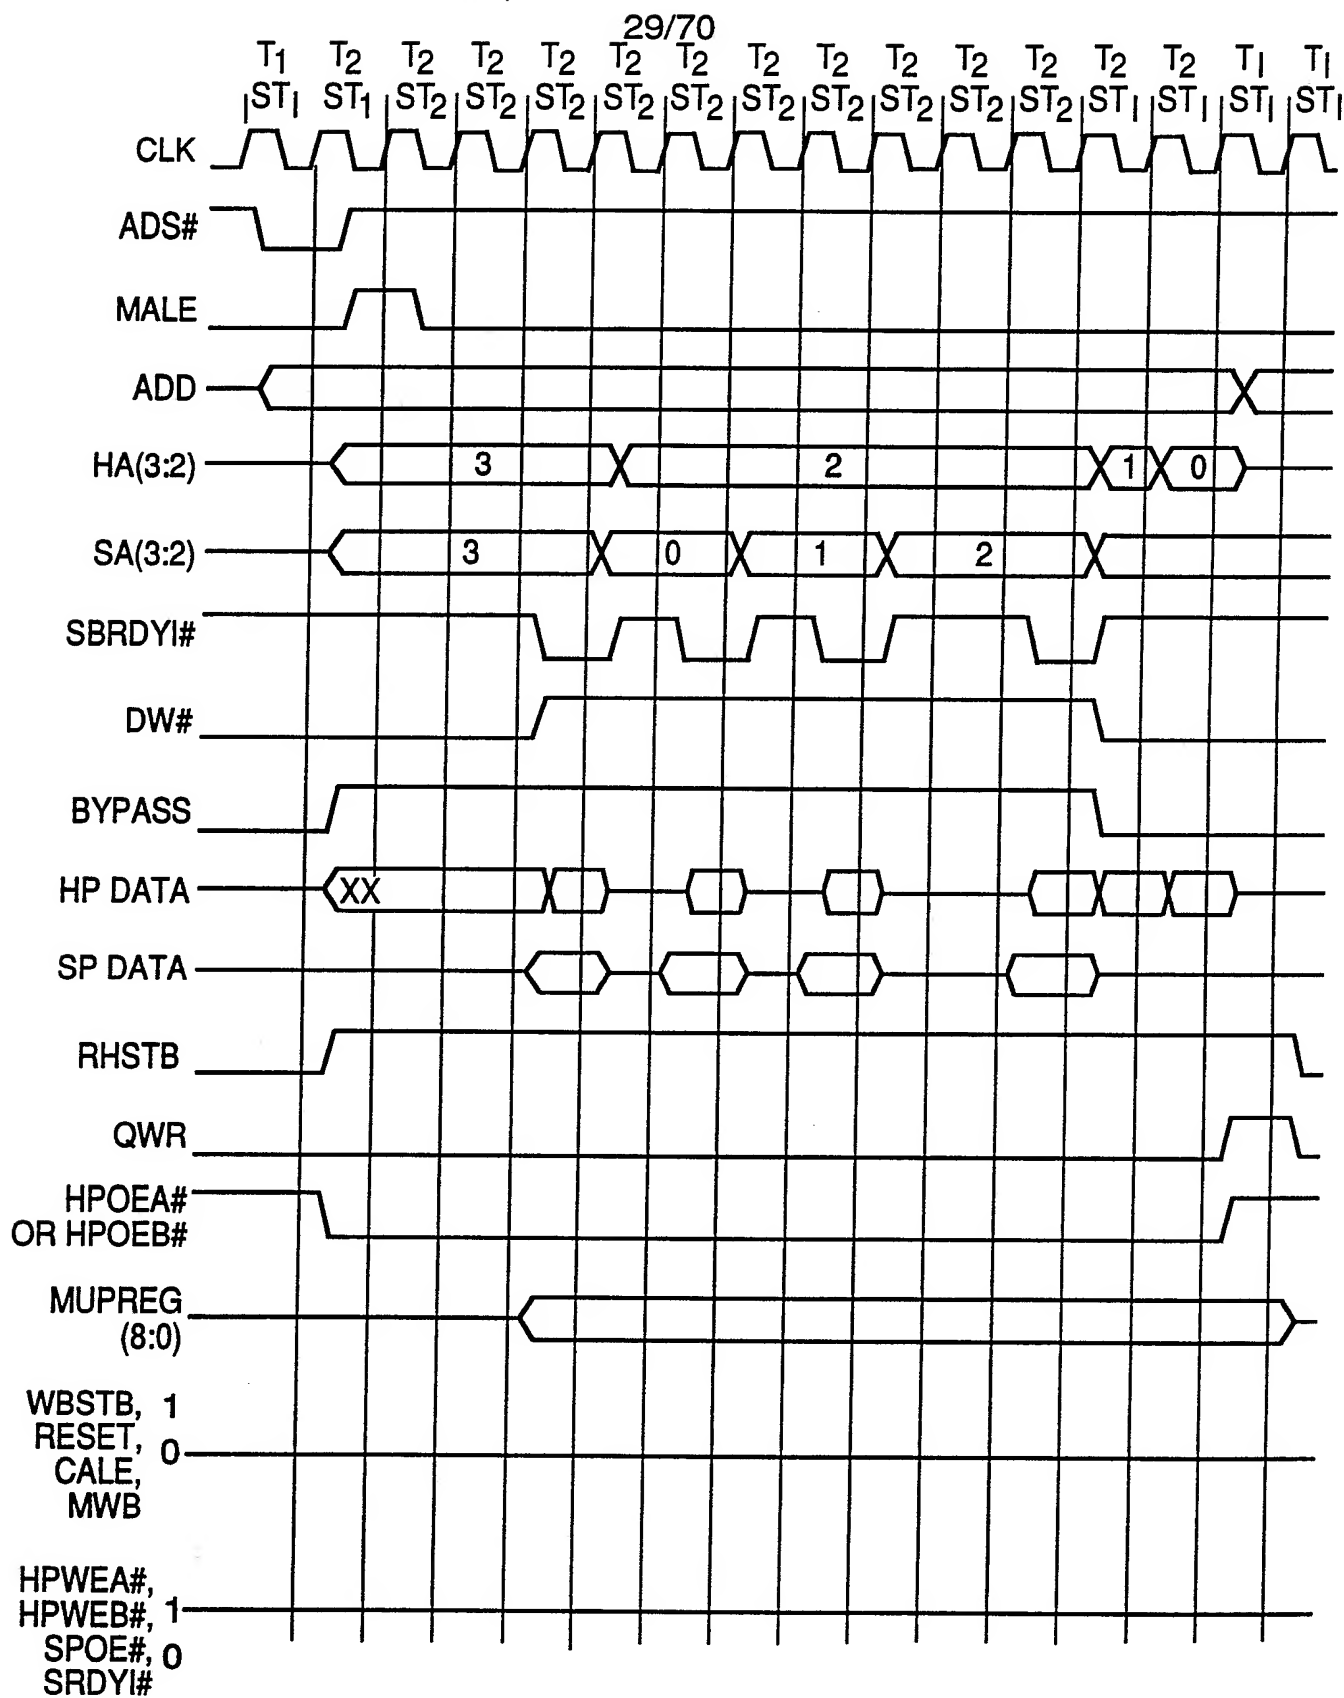



FIG. 24

30/70

03	02	01	00
48	07	66	2c
Valid & Dirty	Non-Valid	Non-Valid	Non-Valid

FIG. 25



32/70




FIG. 27

33/70




FIG. 28

34/70




FIG. 29

35/70

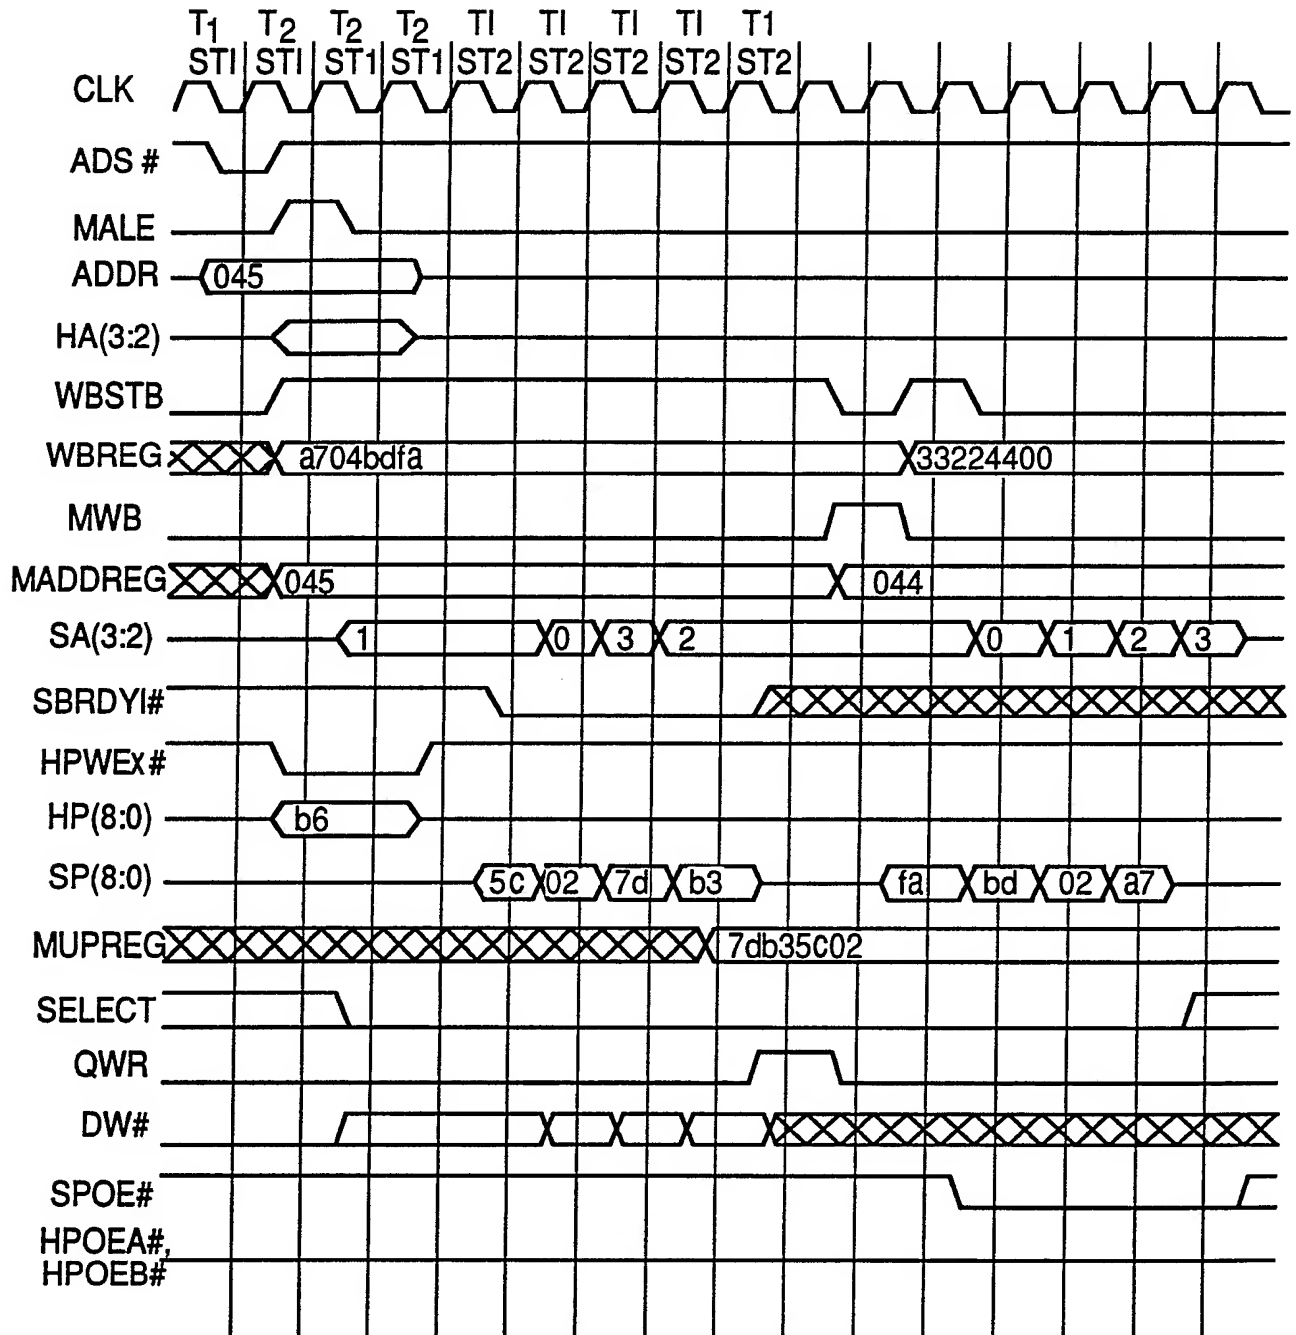


FIG. 30




FIG. 31

37/70




FIG. 32

KEY TO
FIGURE 33

FIG. 33(1) | FIG. 33(2)




FIG. 33(1)




FIG. 33(2)

40/70

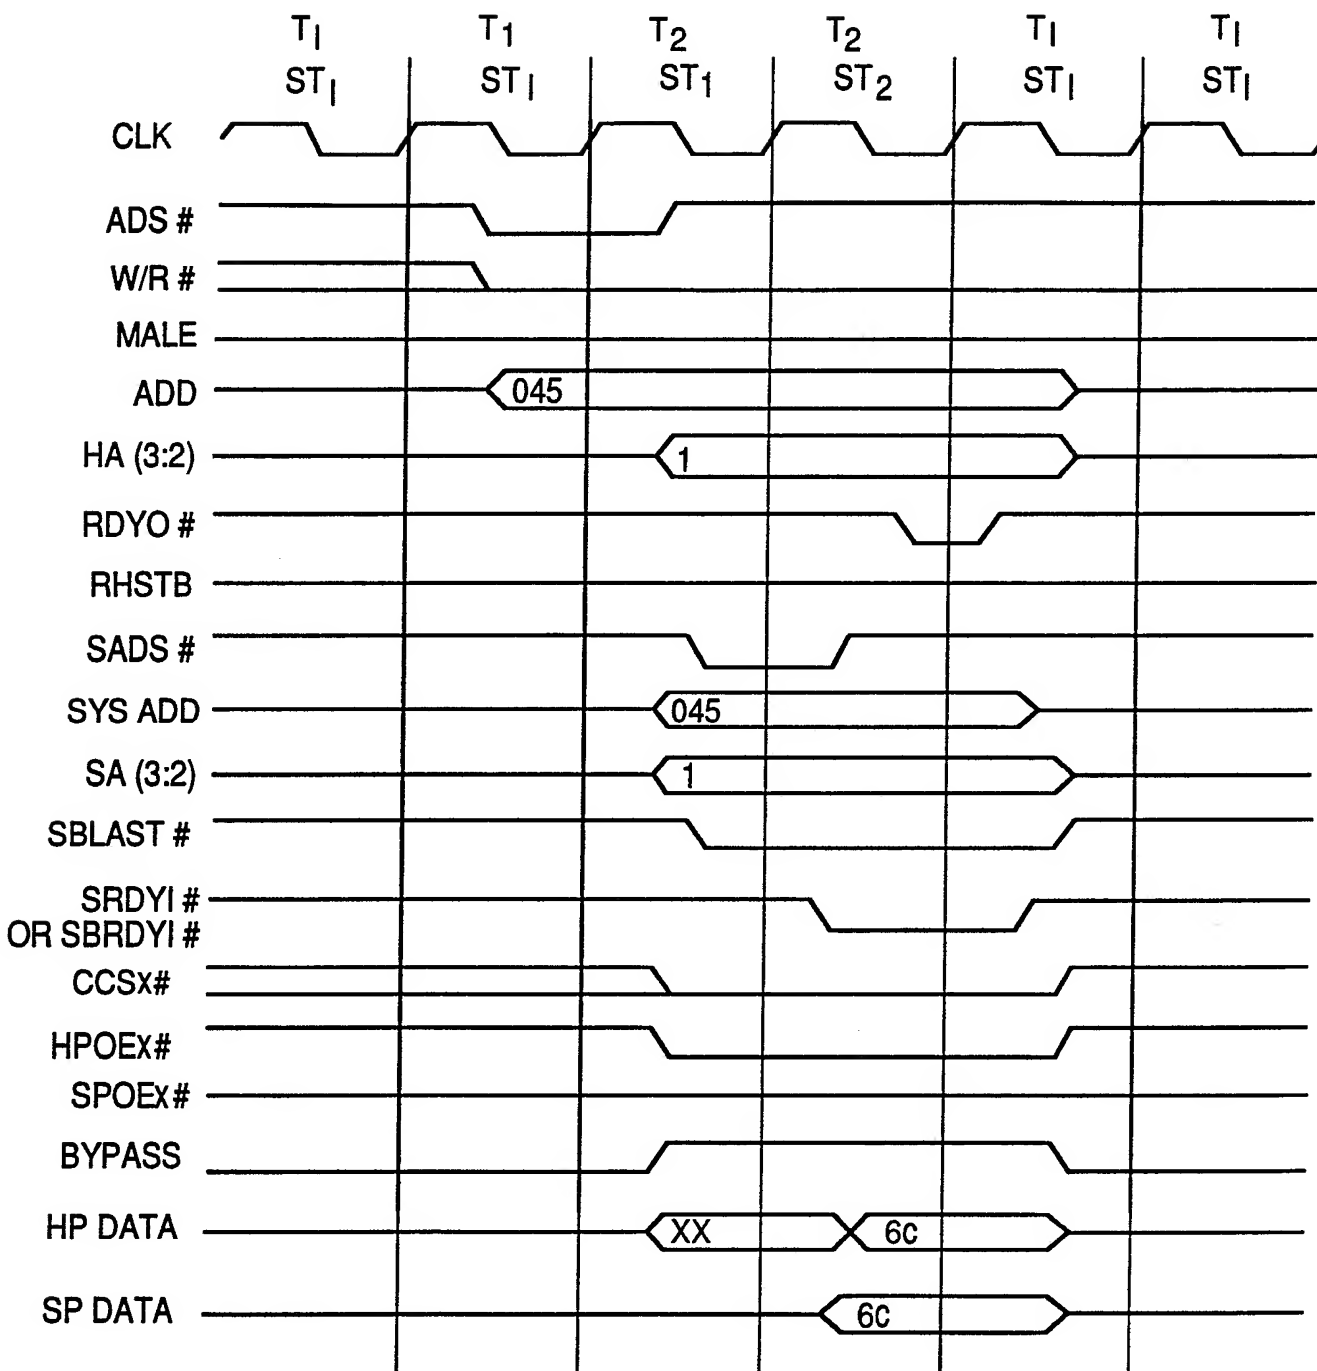


FIG. 34

41/70

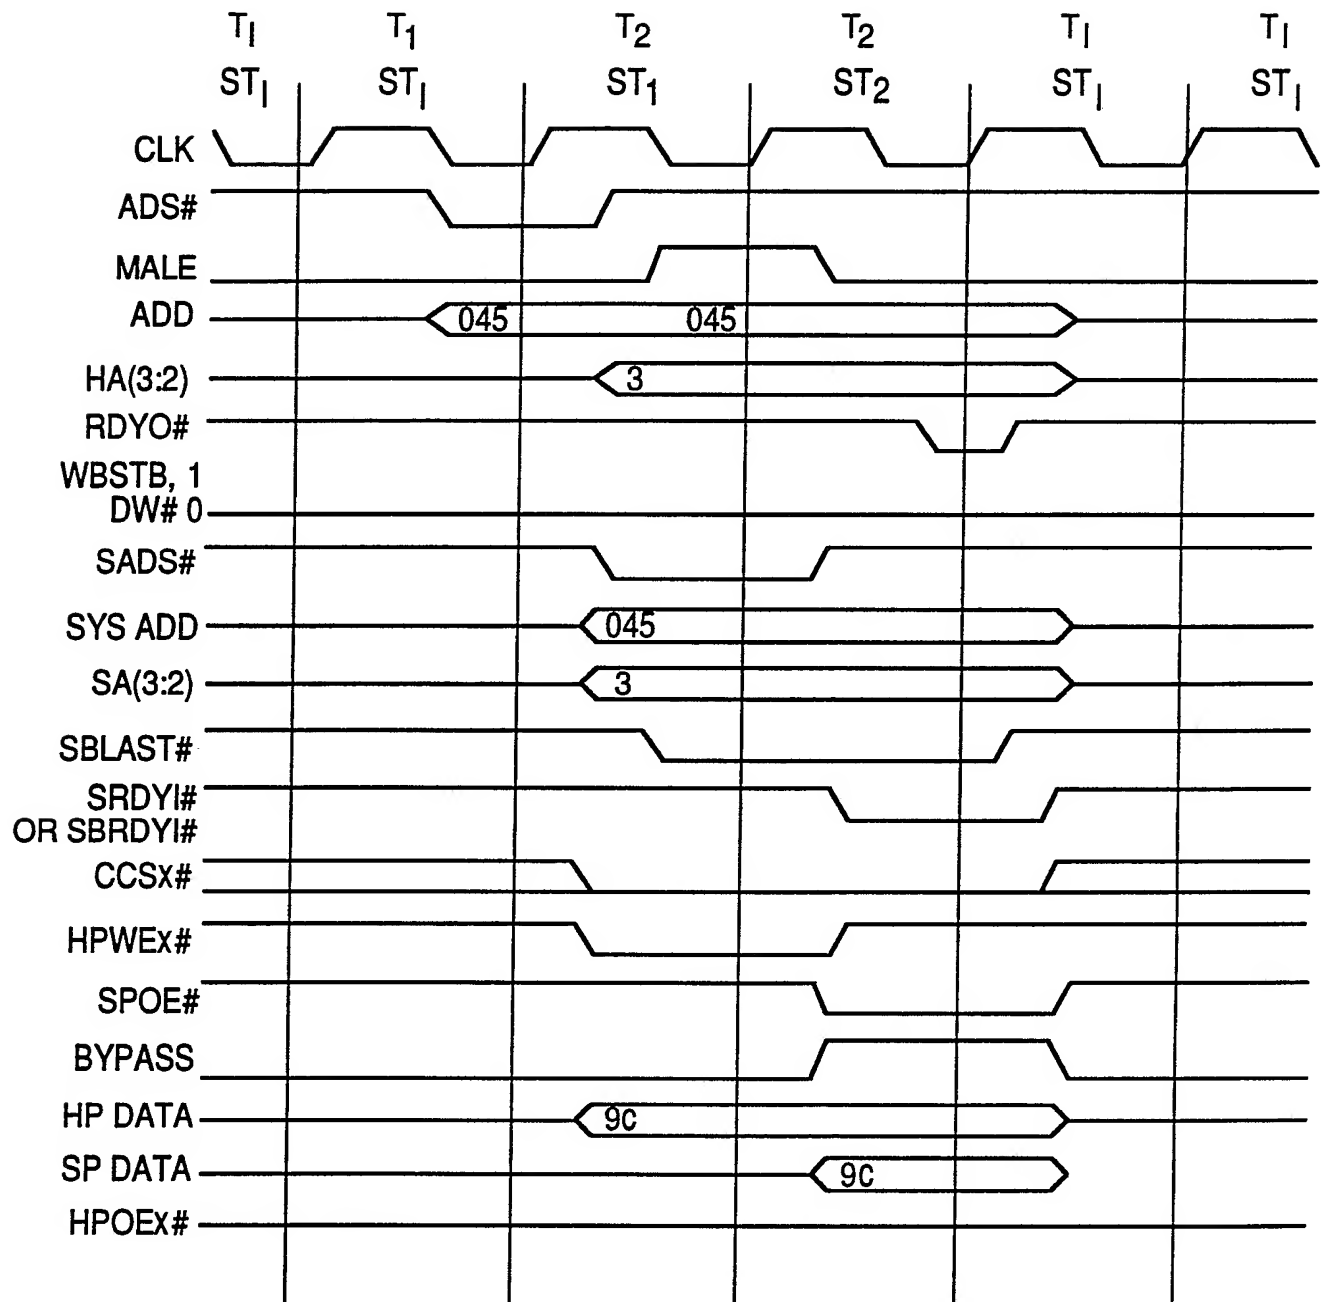


FIG. 35

42/70

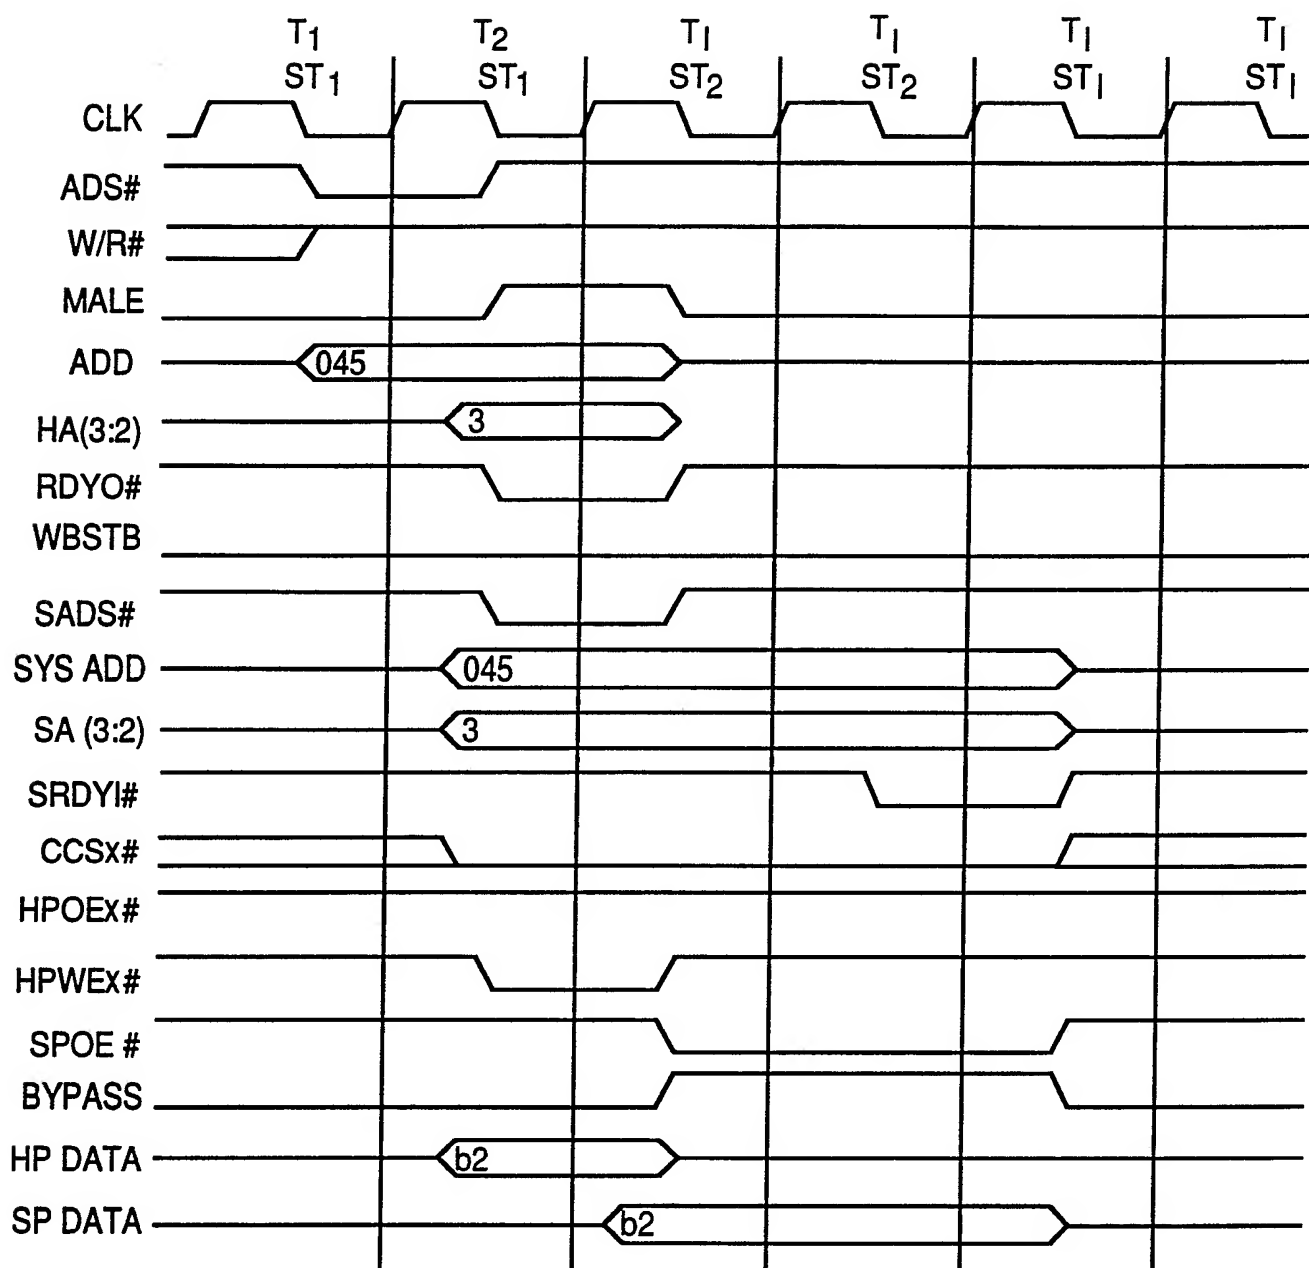


FIG. 36

43/70

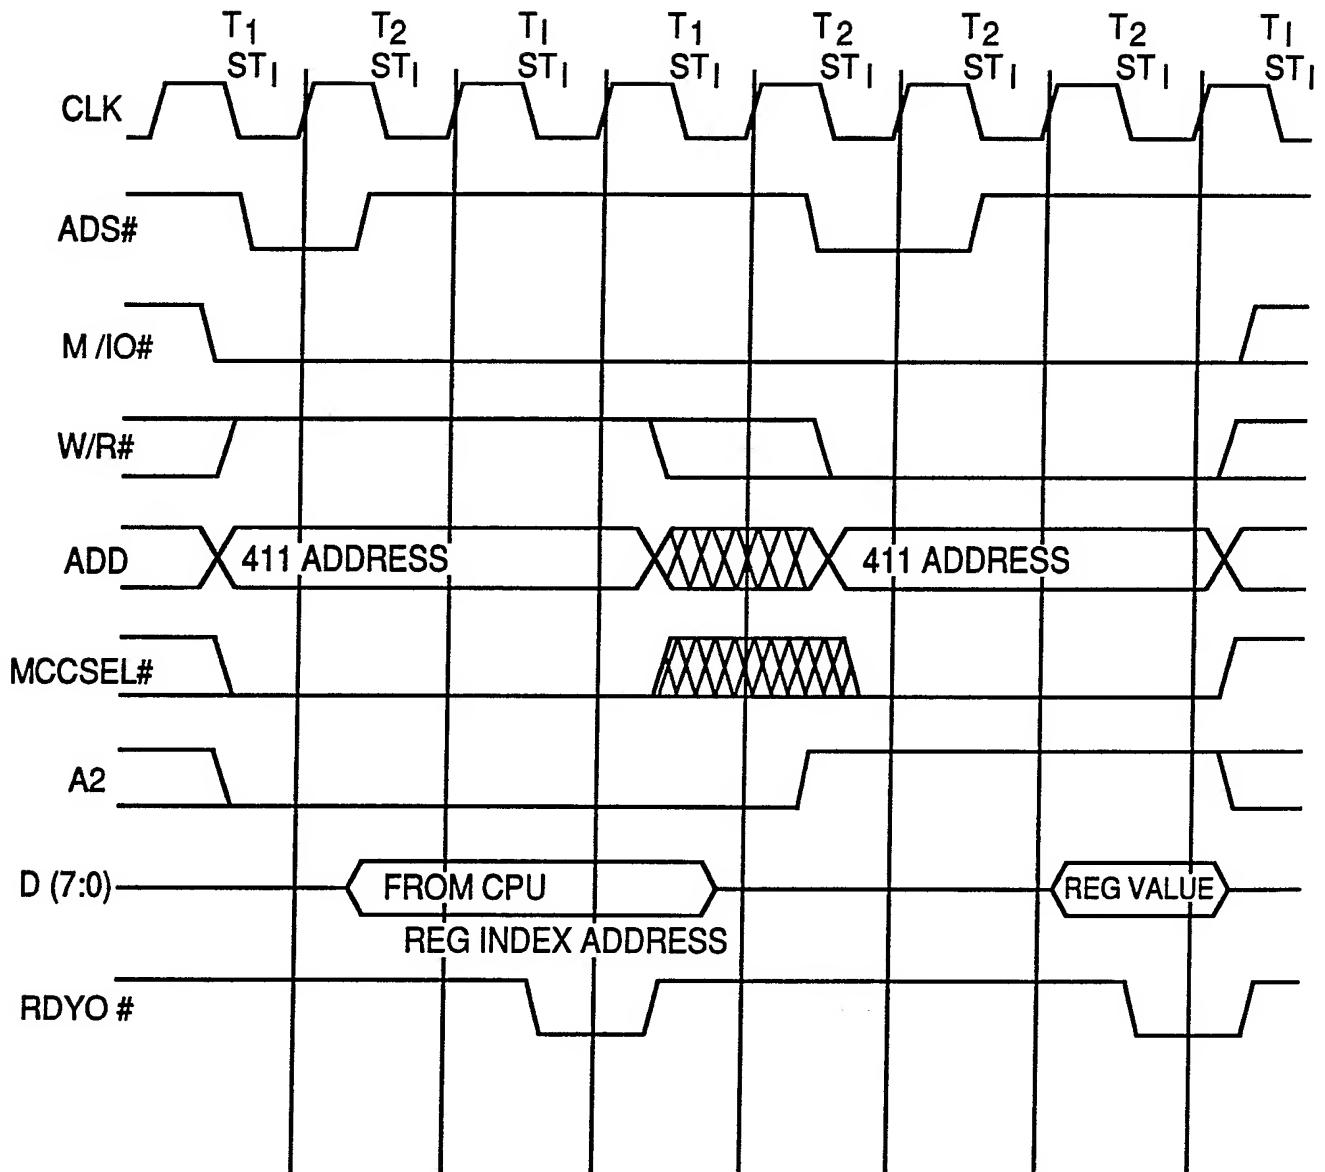


FIG. 37

44/70




FIG. 38

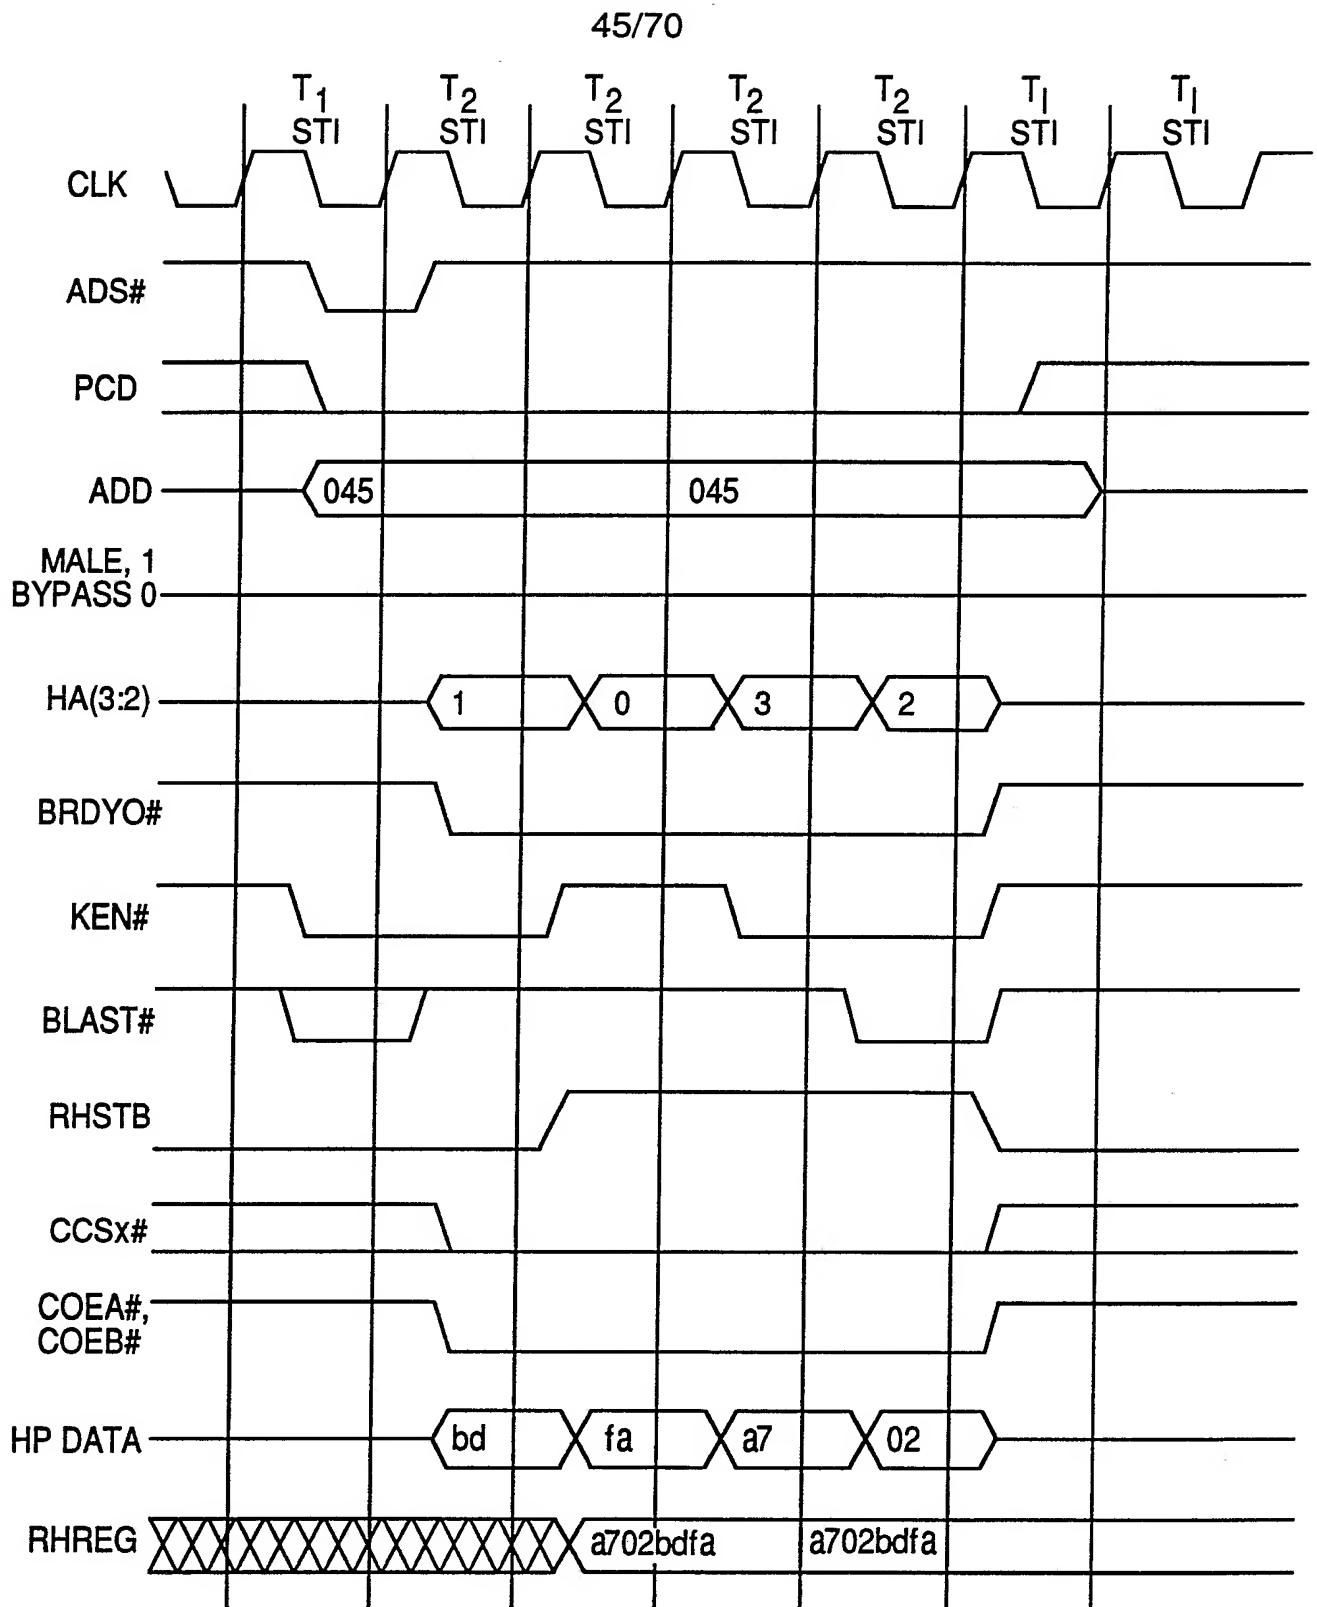


FIG. 39

46/70

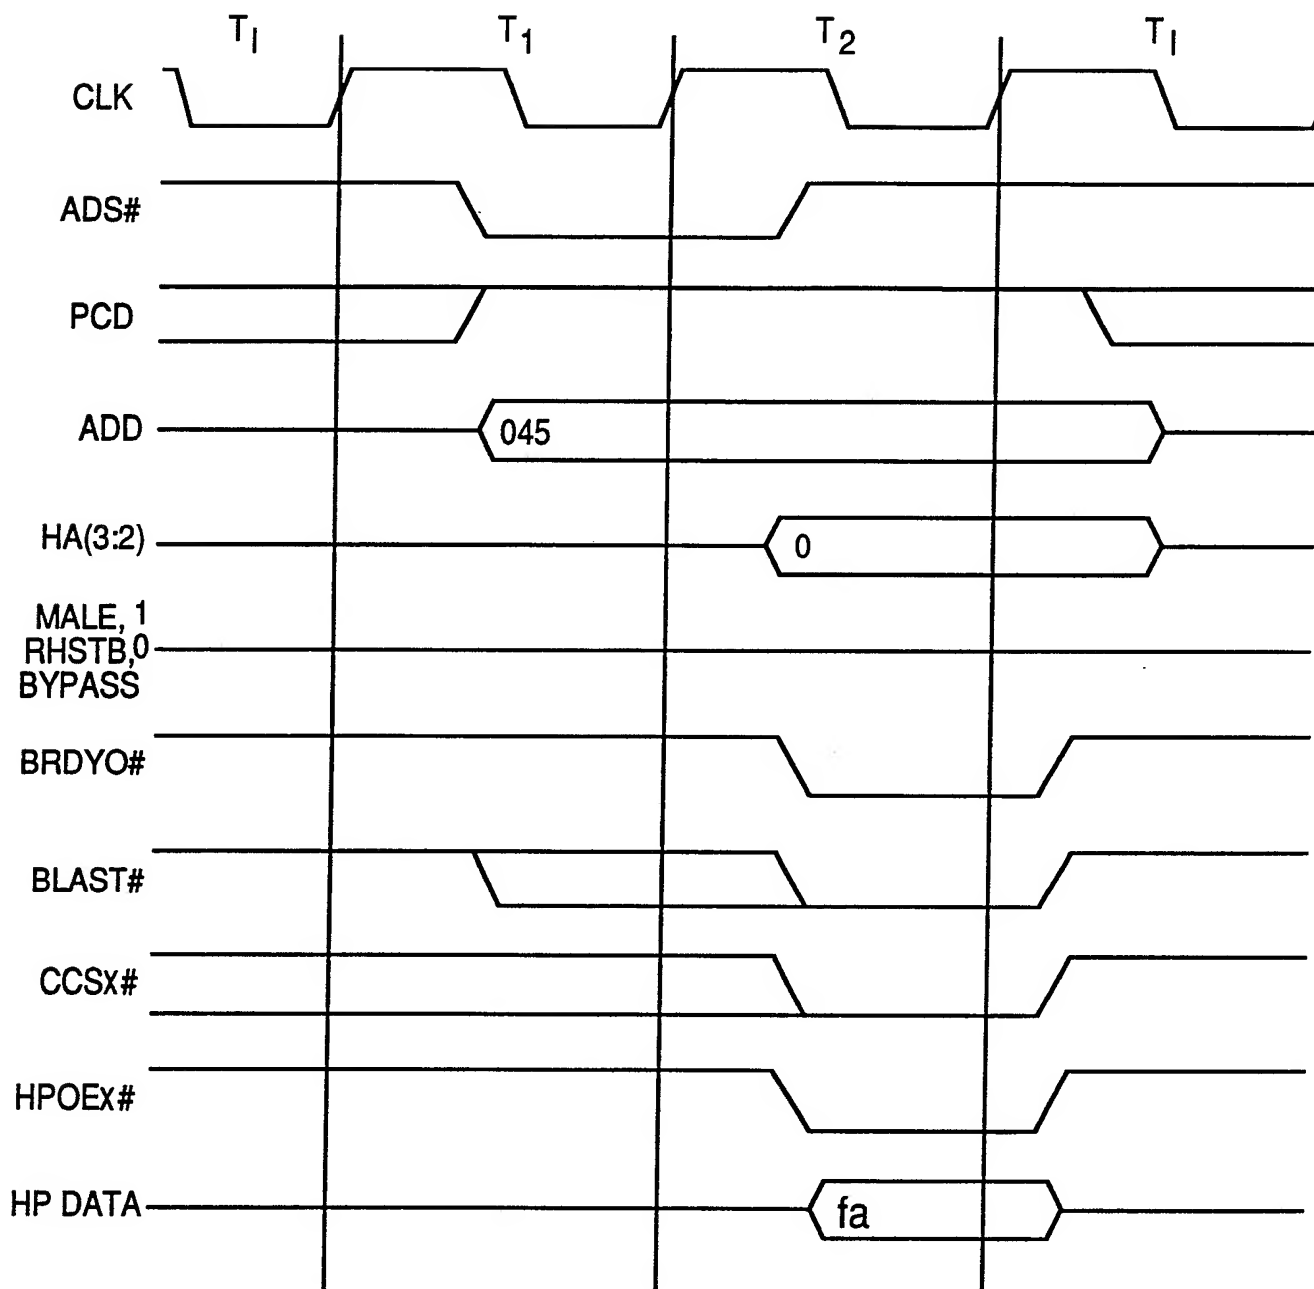


FIG. 40

47/70




FIG. 41

48/70




FIG. 42

49/70

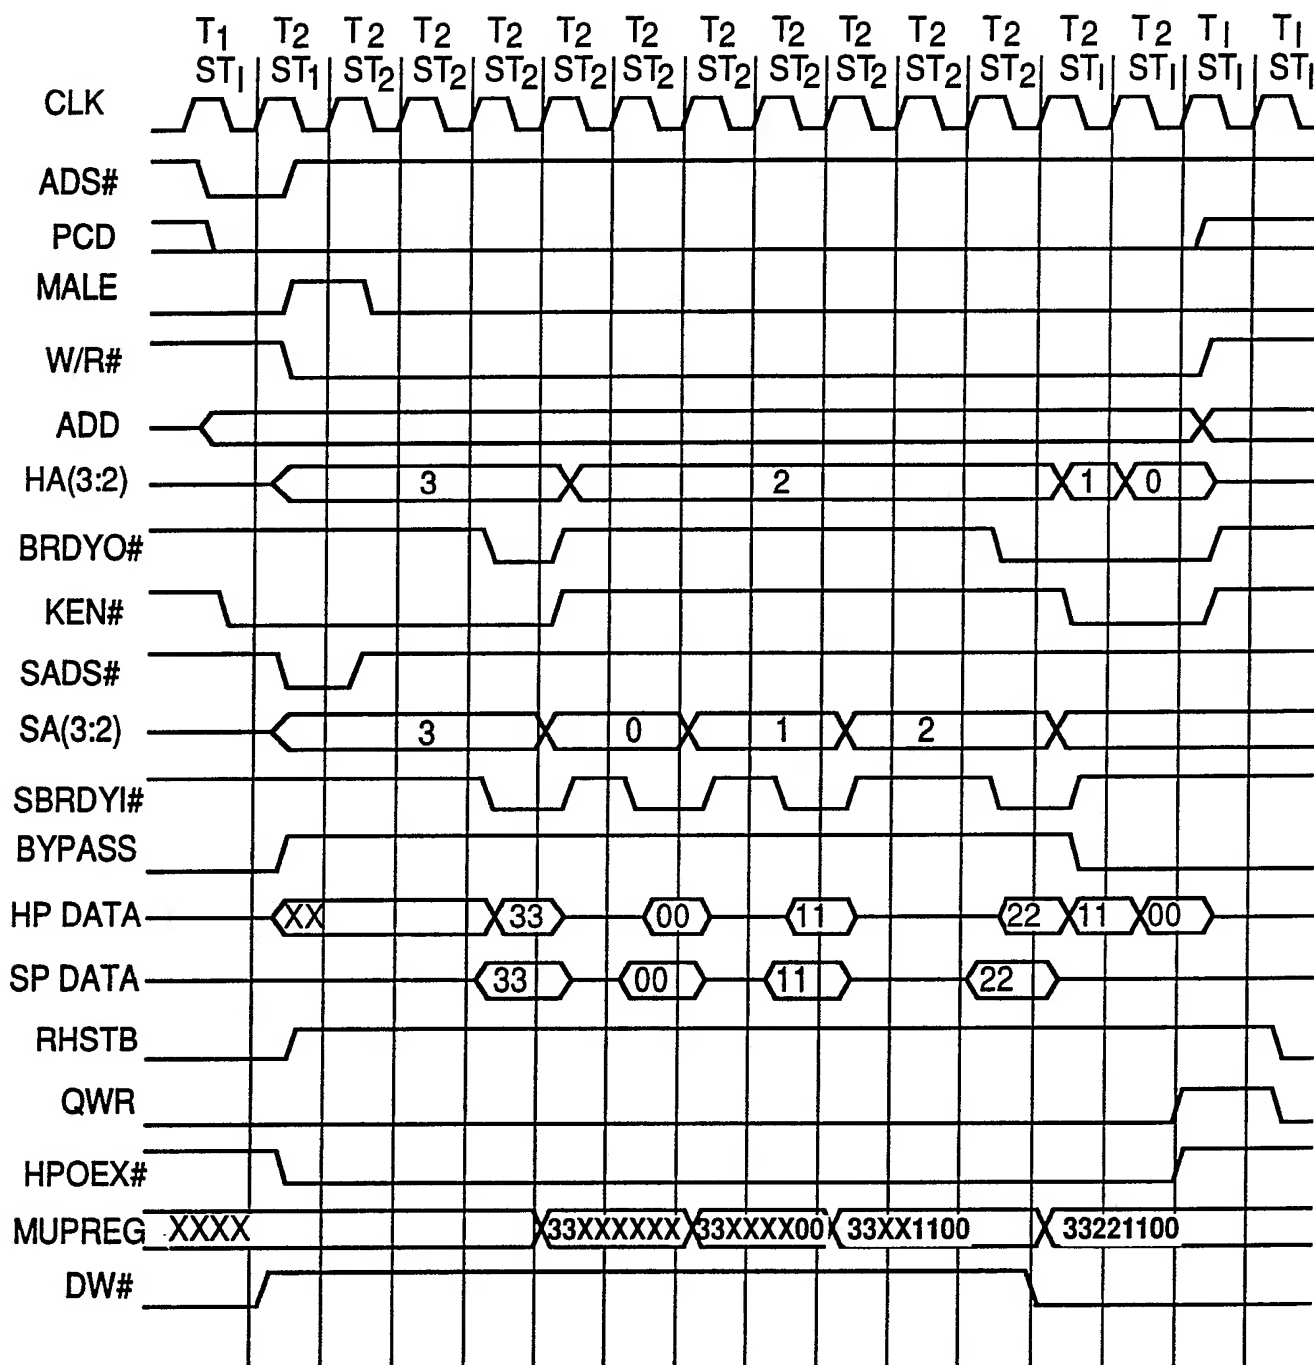


FIG. 43

50/70

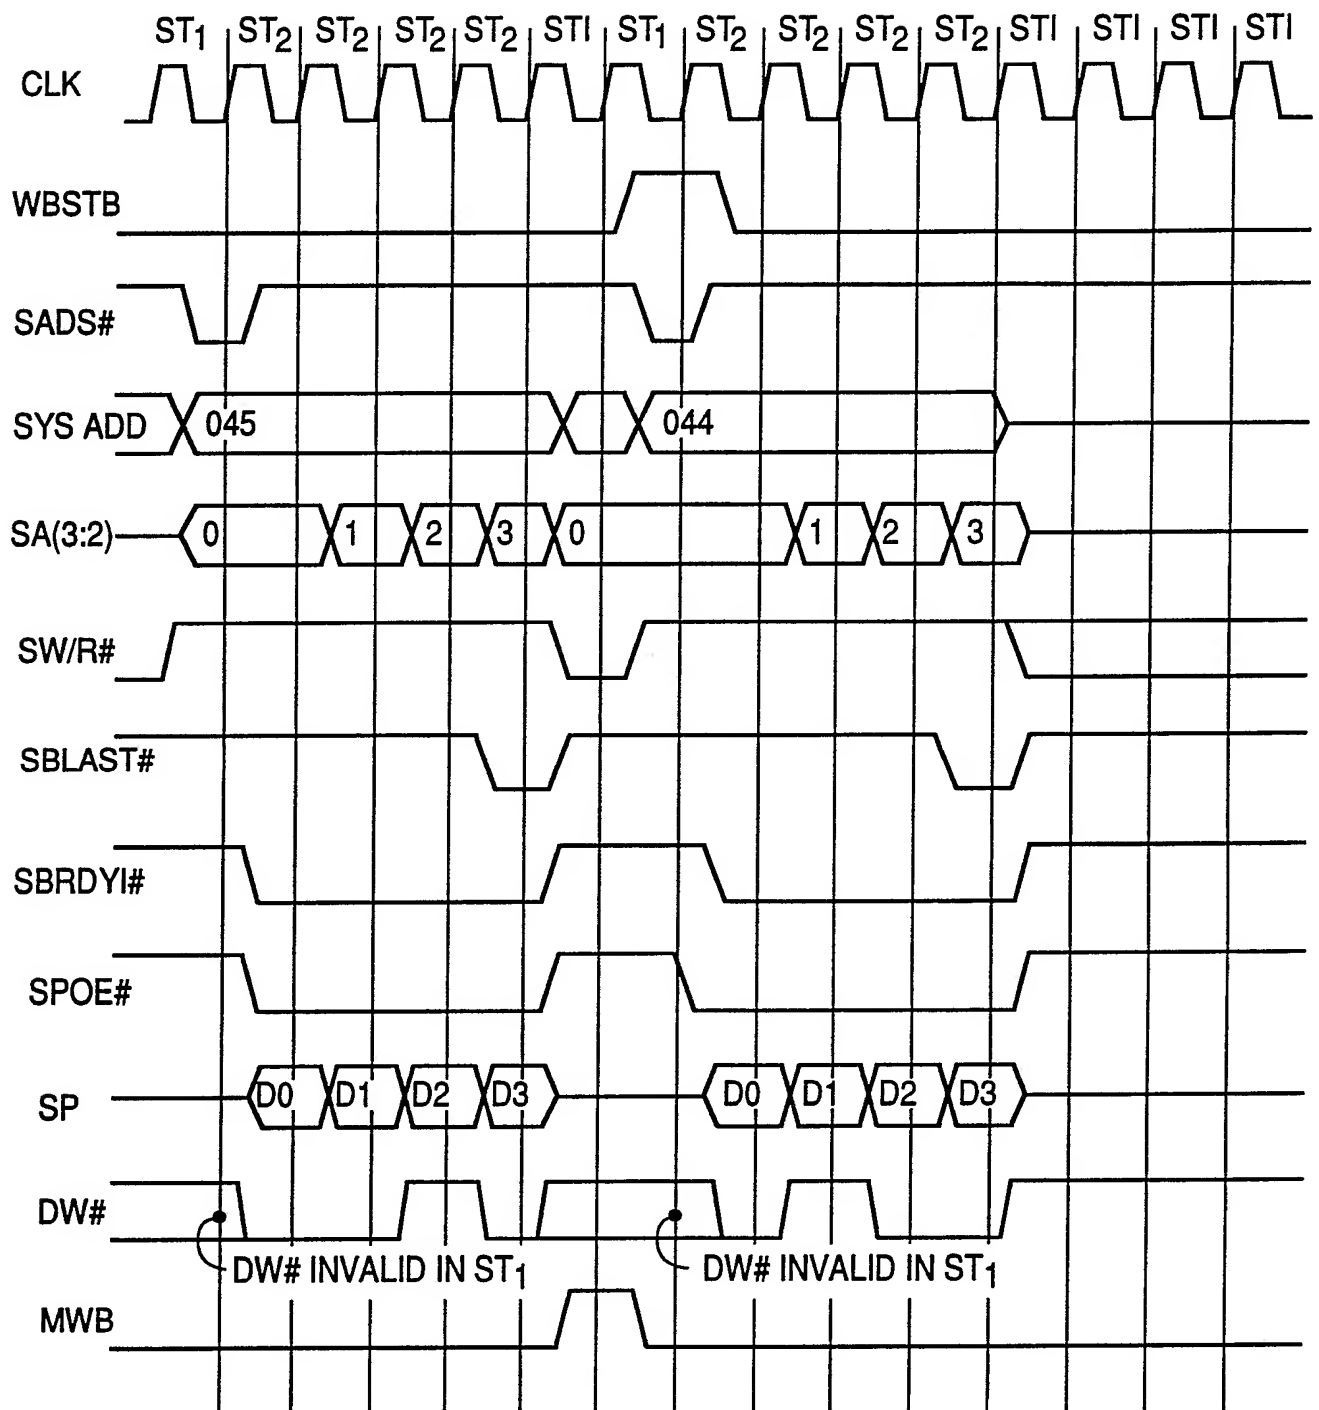


FIG. 44

51/70

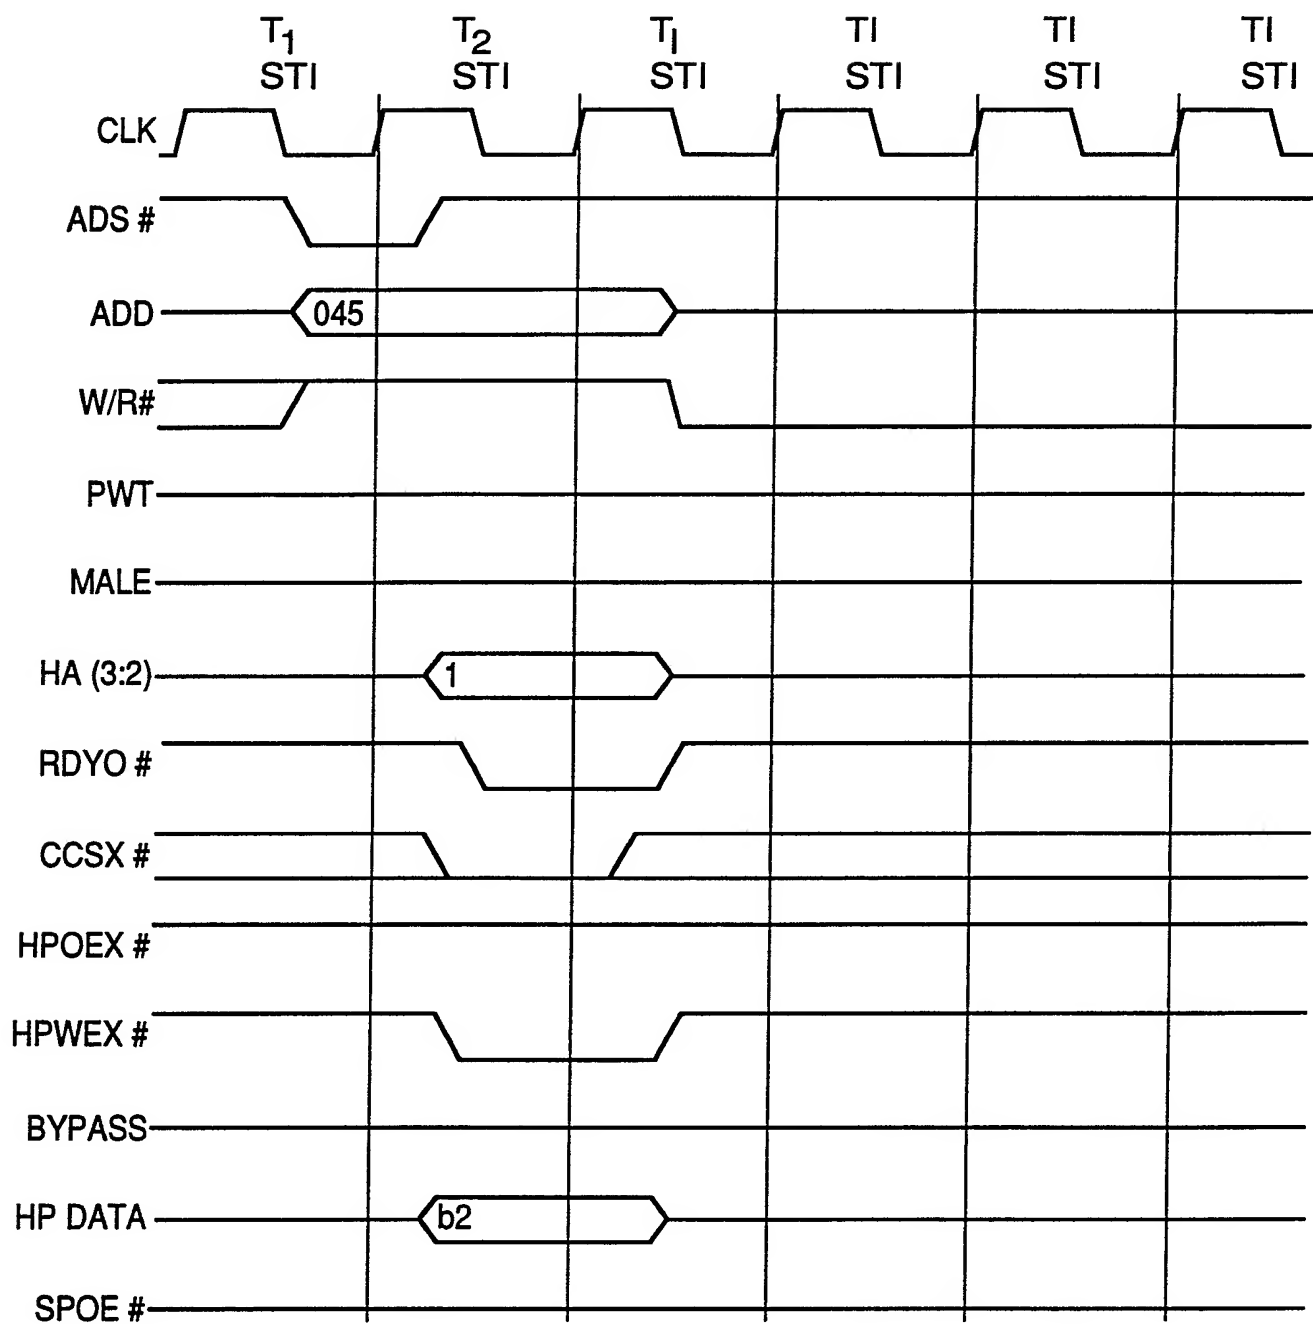


FIG. 45




FIG. 46

53/70

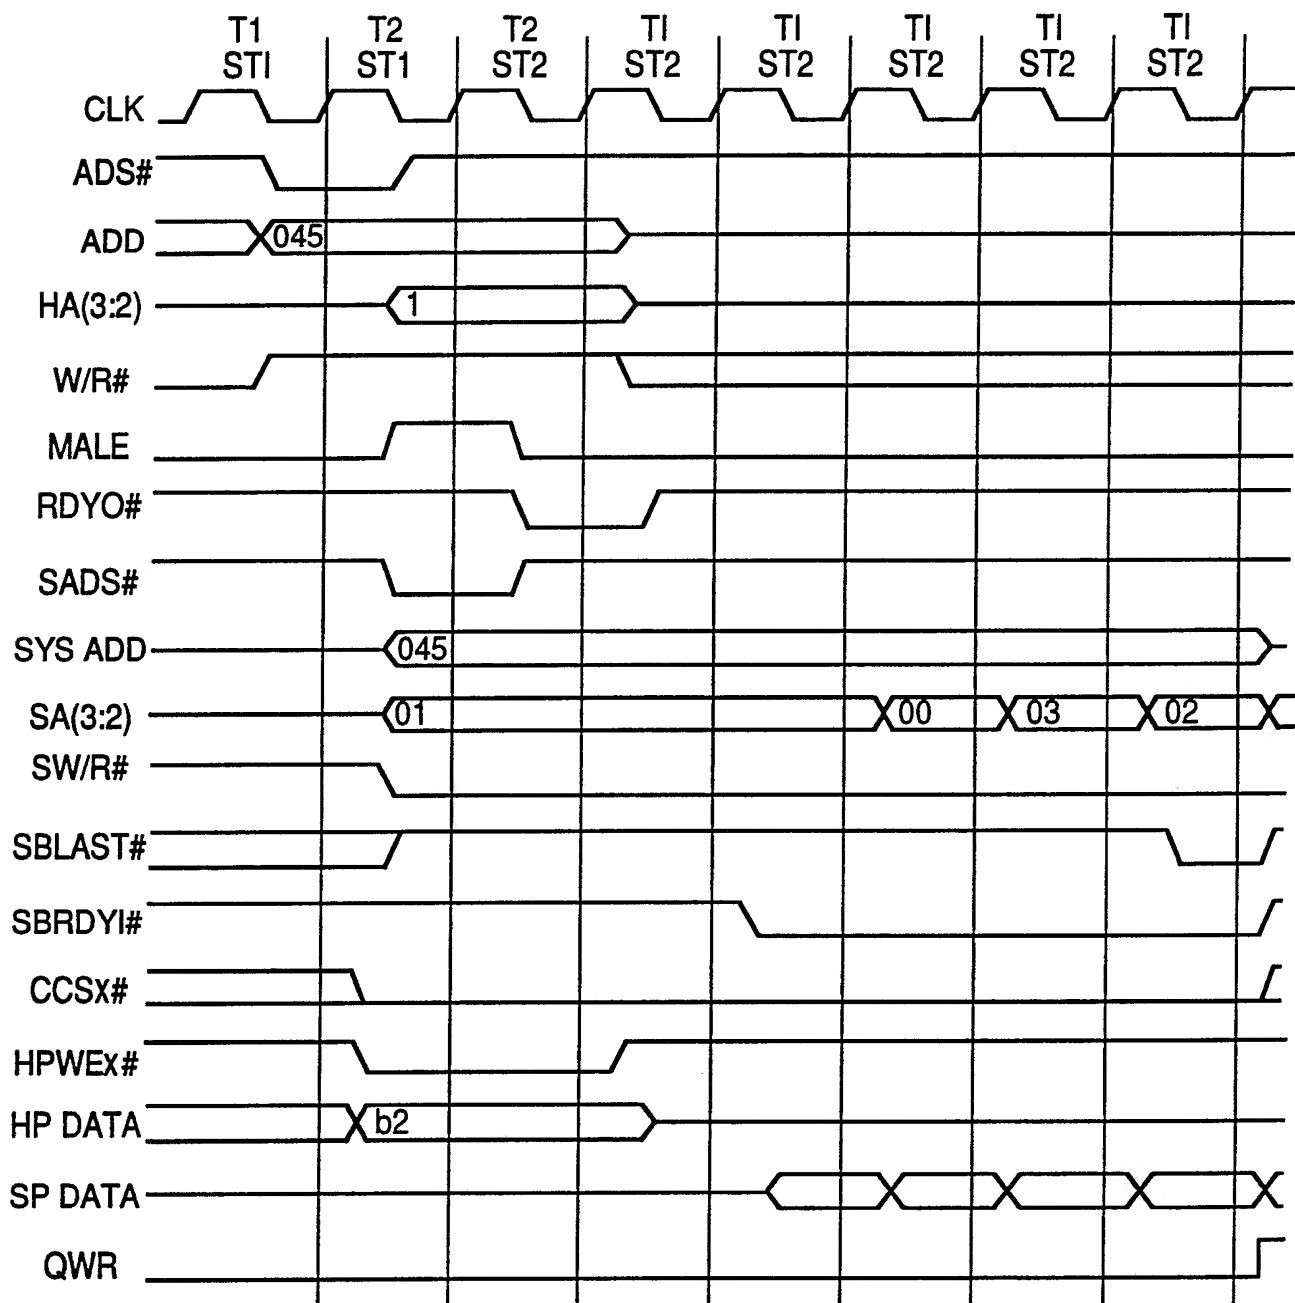


FIG. 47

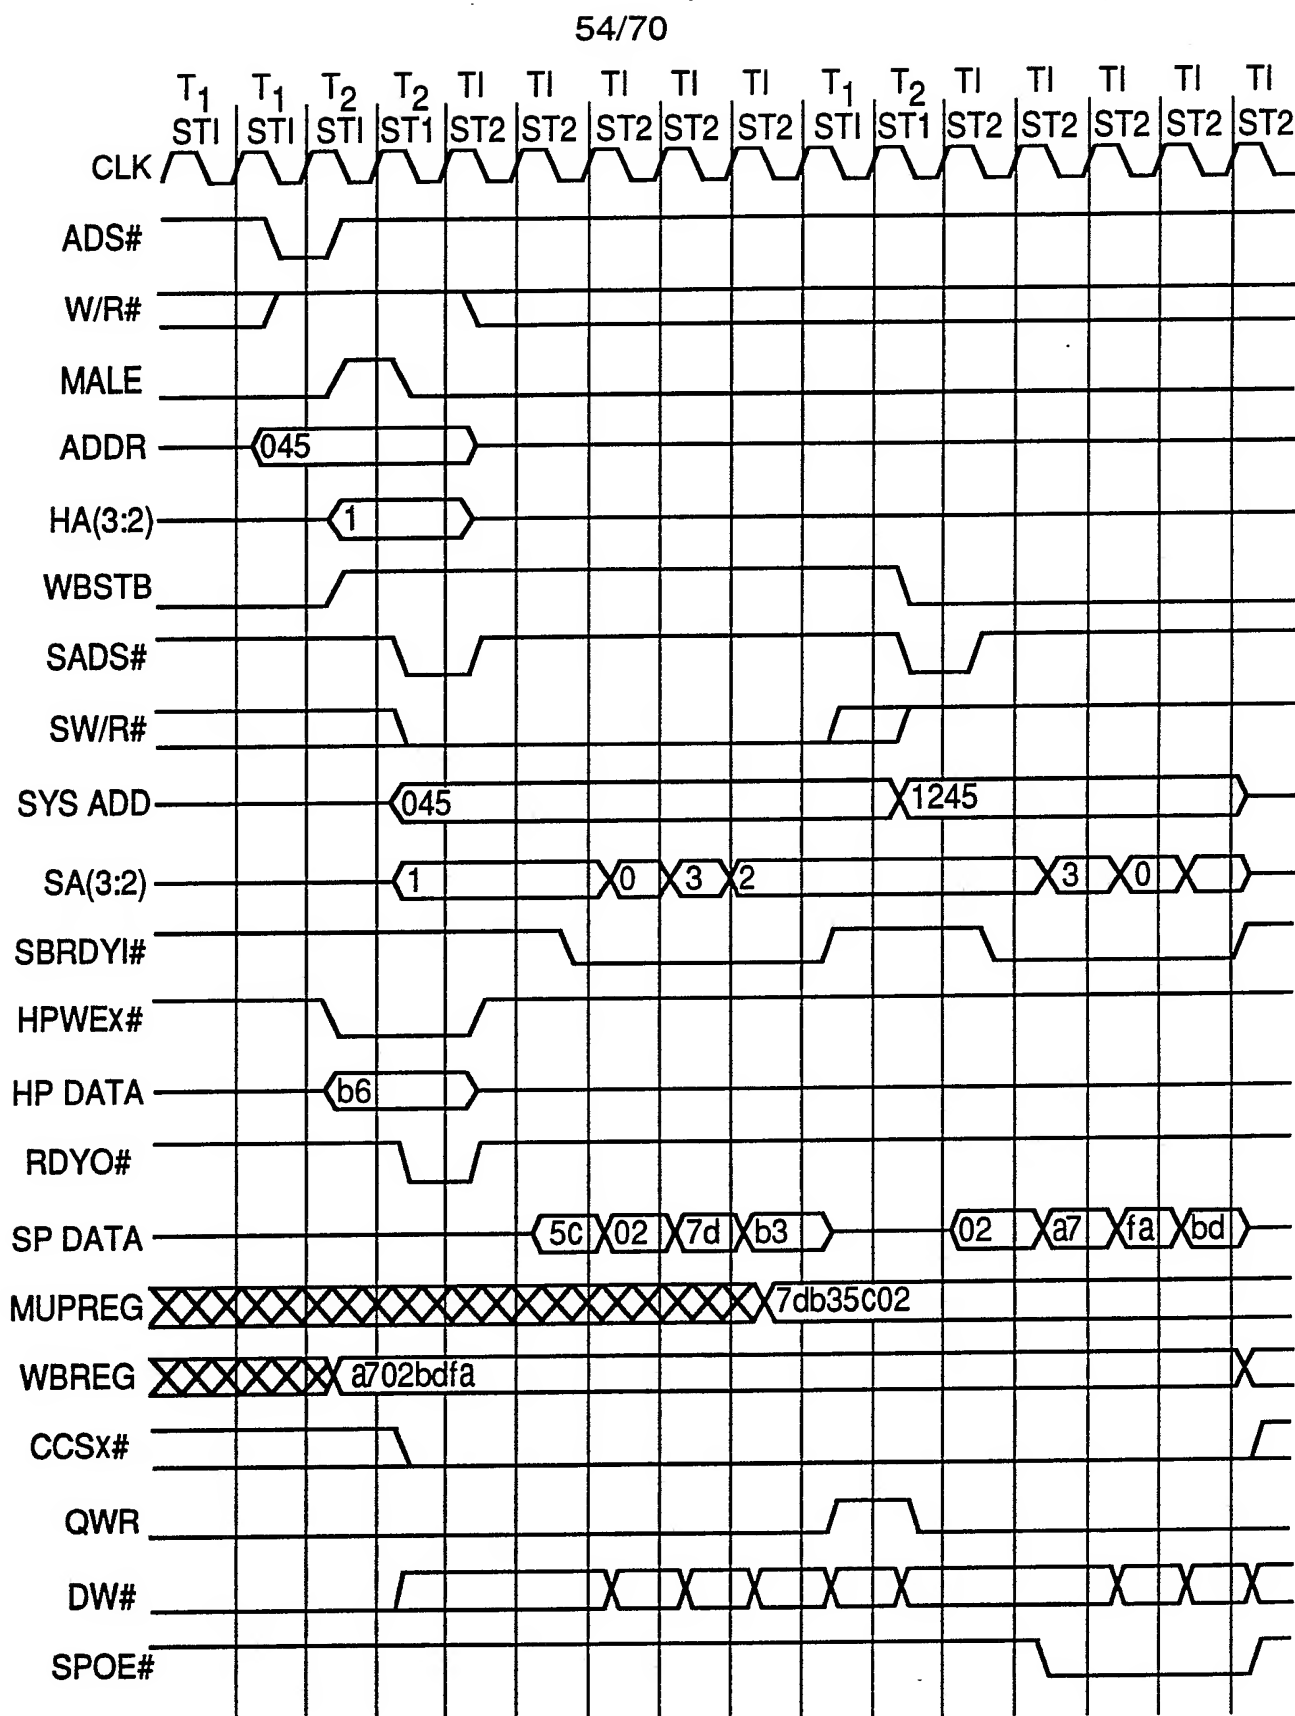


FIG. 48

55/70




FIG. 49

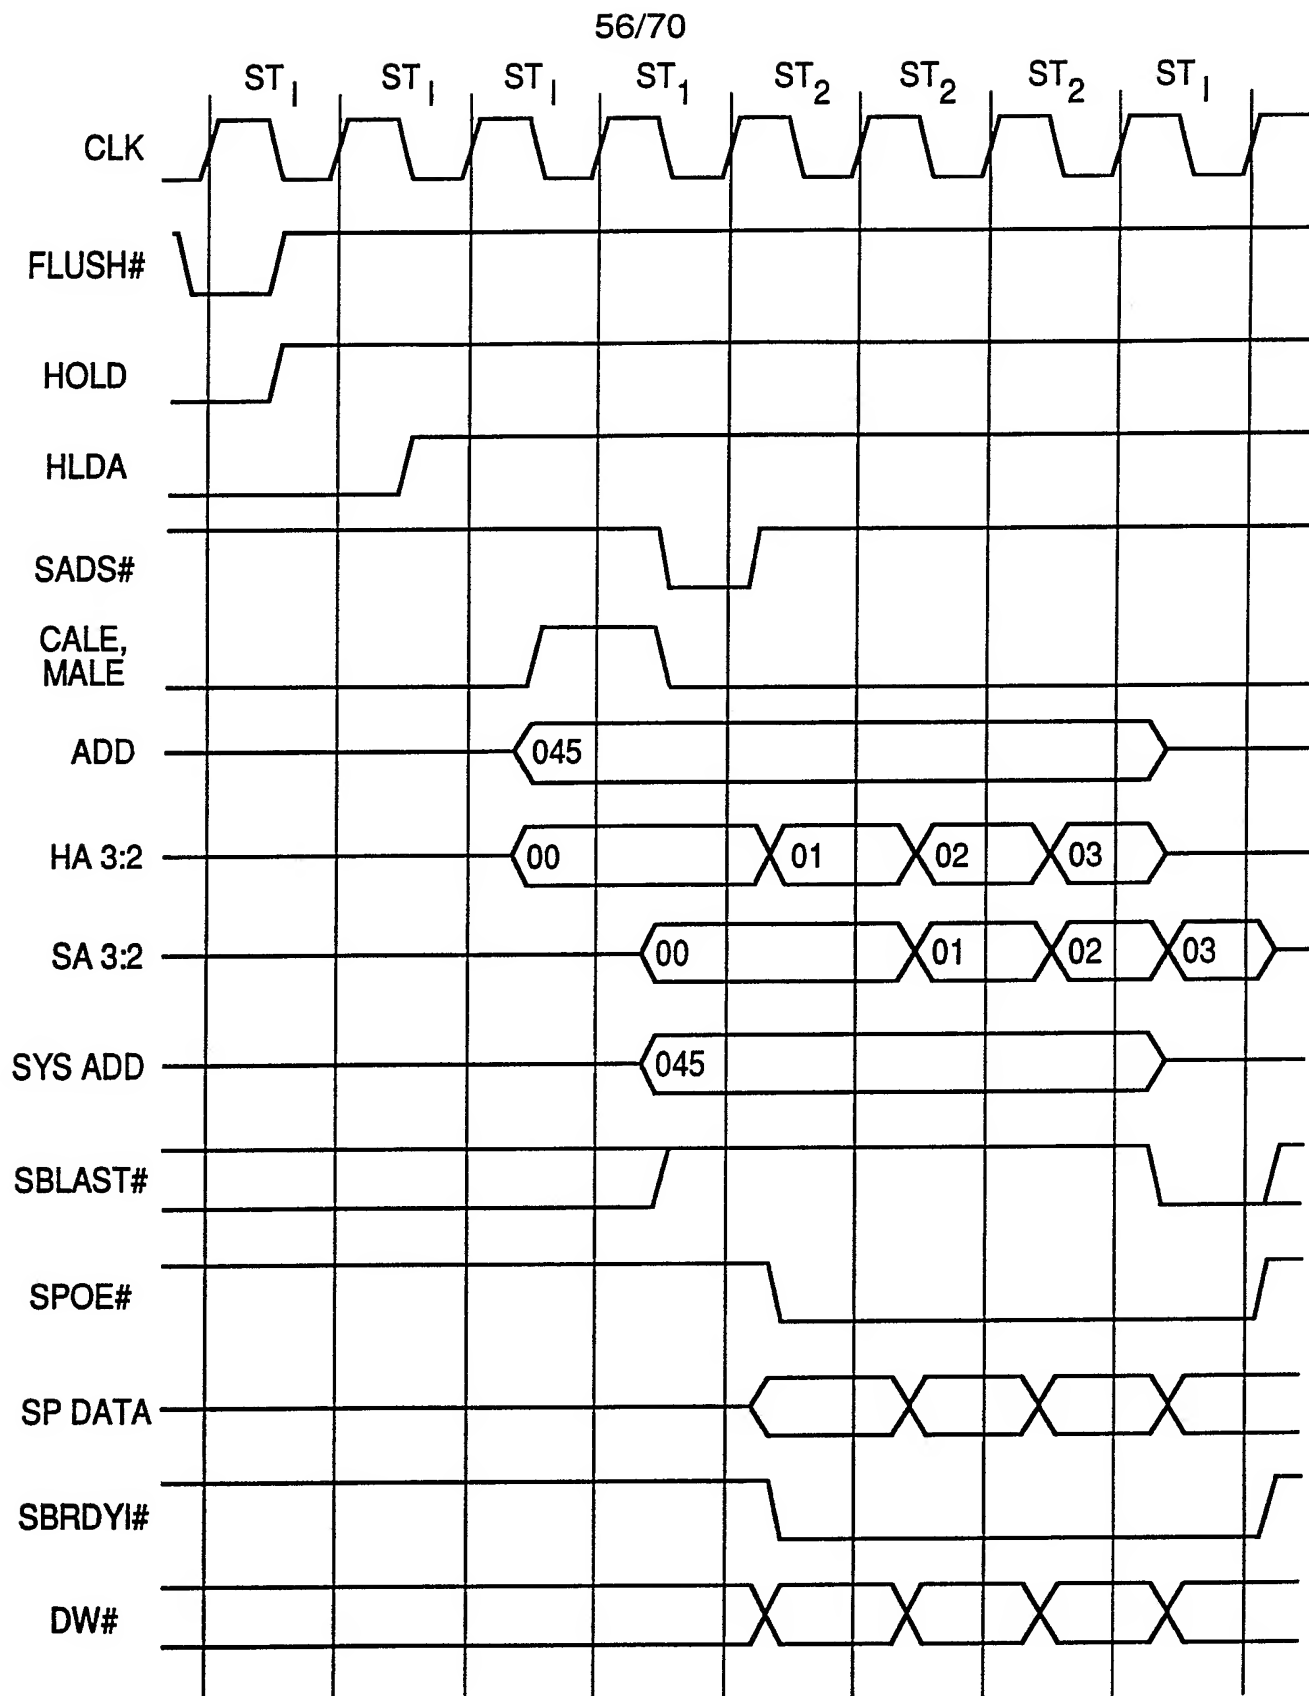



FIG. 50

57/70

FIG. 51



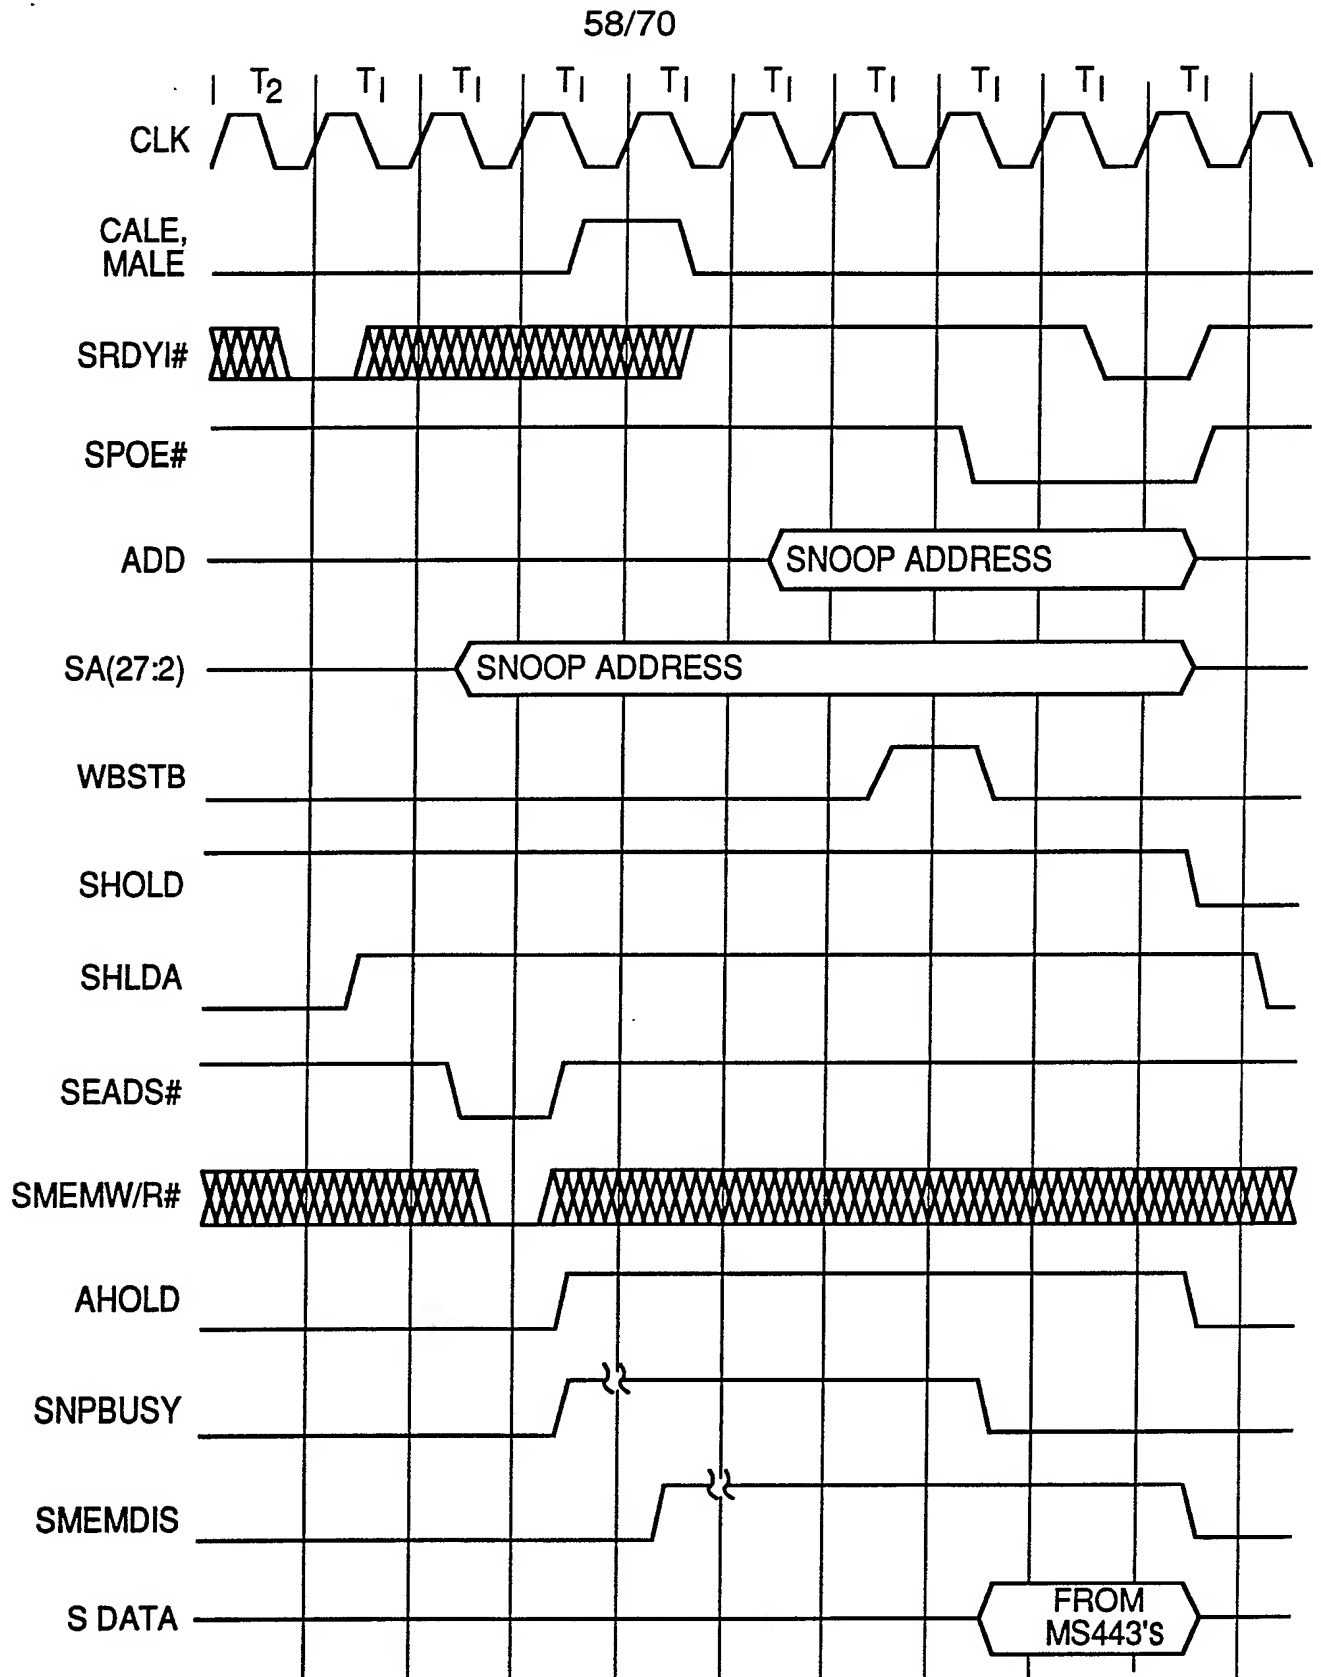


FIG. 52

59/70

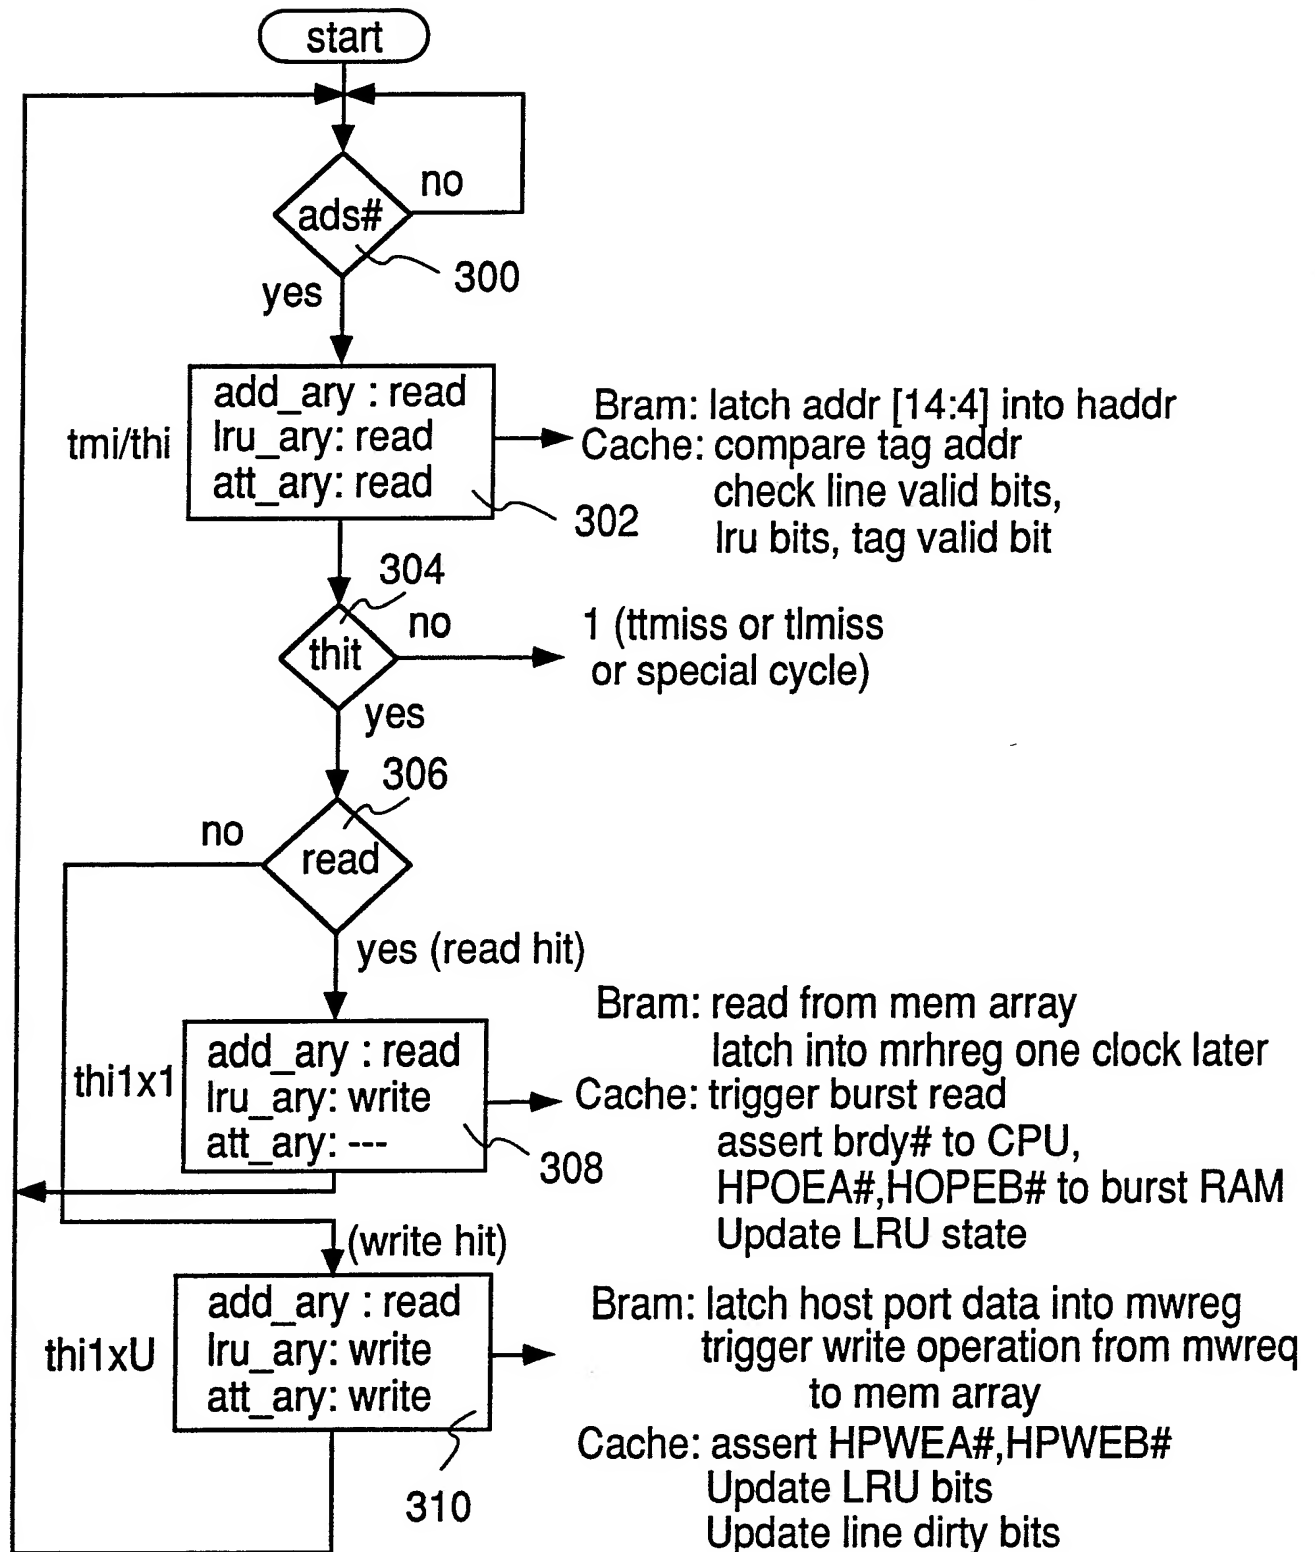
441/443 Sequencing

FIG. 53

Read Tag Miss




FIG. 54(1)

61/70





FIG. 54(2)

62/70



KEY TO
FIGURE 54
FIG. 54(1)
FIG. 54(2)
FIG. 54(3)

FIG. 54(3)

Write Tag Miss

- p.s. 1. tc4d : current line contain dirty data
- 2. tn4d : next line contain dirty data




FIG. 55(1)

64/70





FIG. 55(2)

65/70



KEY TO

FIGURE 55

FIG. 55(1)

FIG. 55(2)

FIG. 55(3)

FIG. 55(3)

66/70





FIG. 56(1)

67/70






FIG. 56(3)

KEY TO
FIGURE 56

FIG. 56(1)
FIG. 56(2)
FIG. 56(3)

69/70

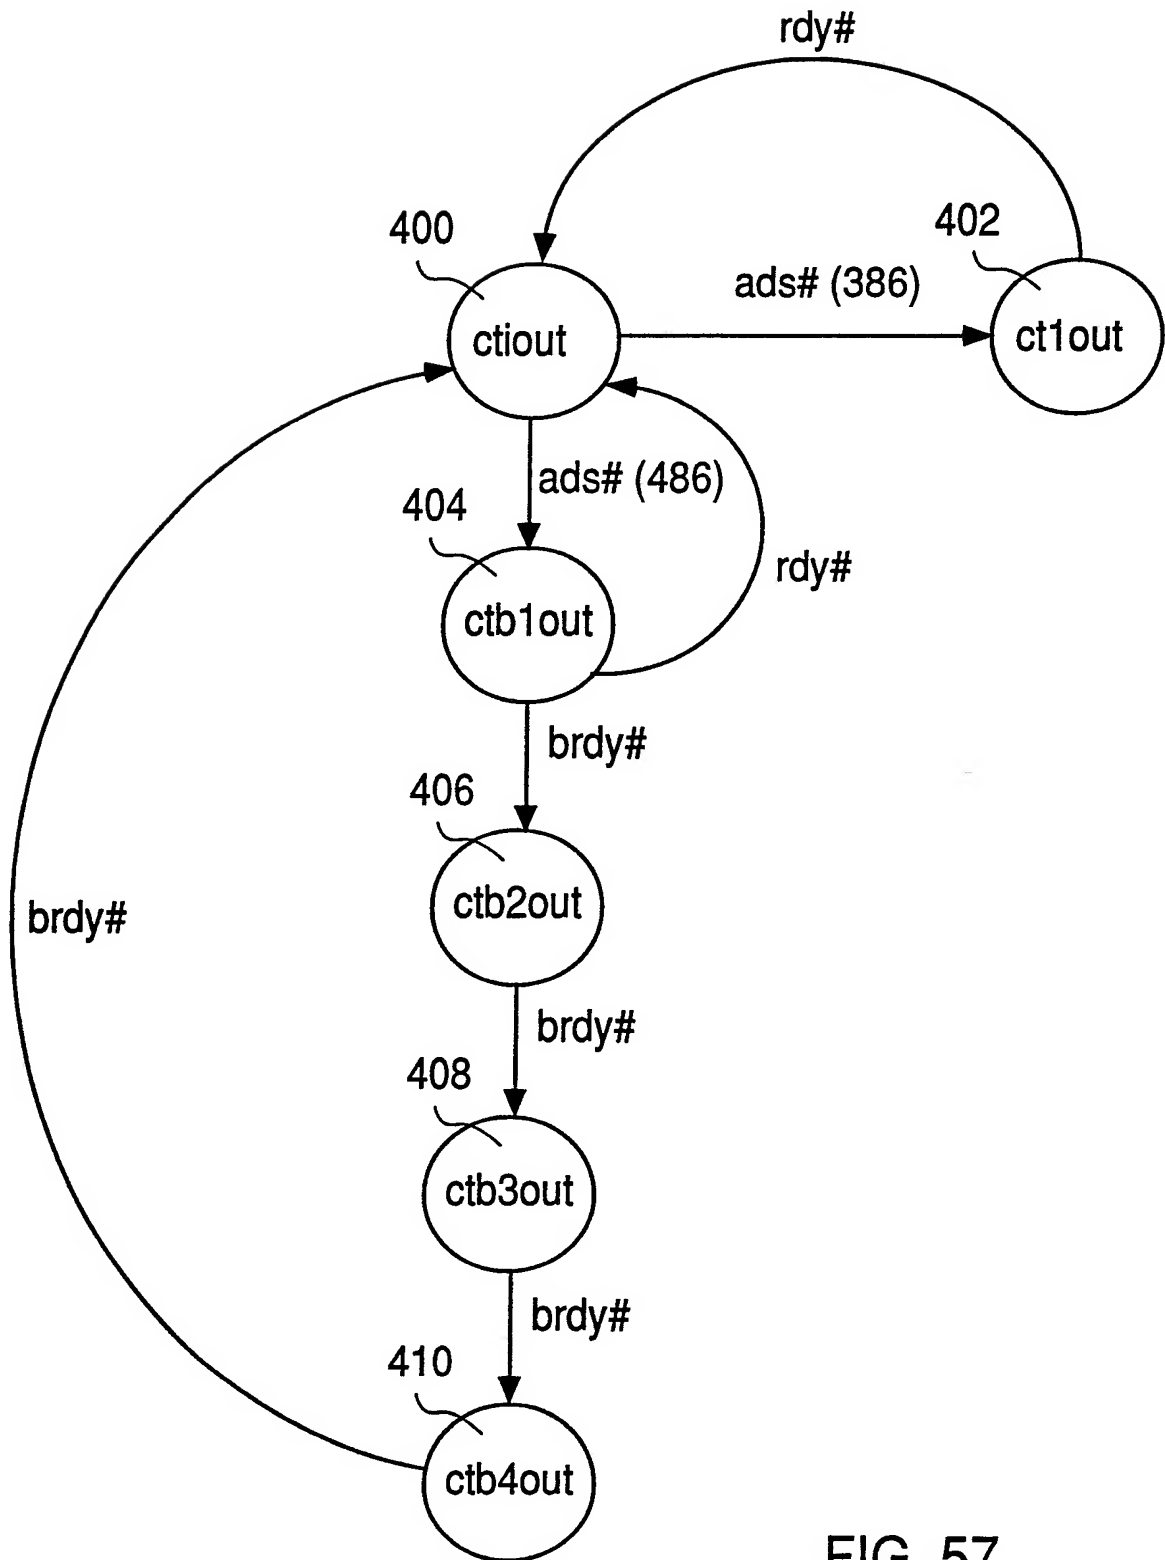


FIG. 57

70/70

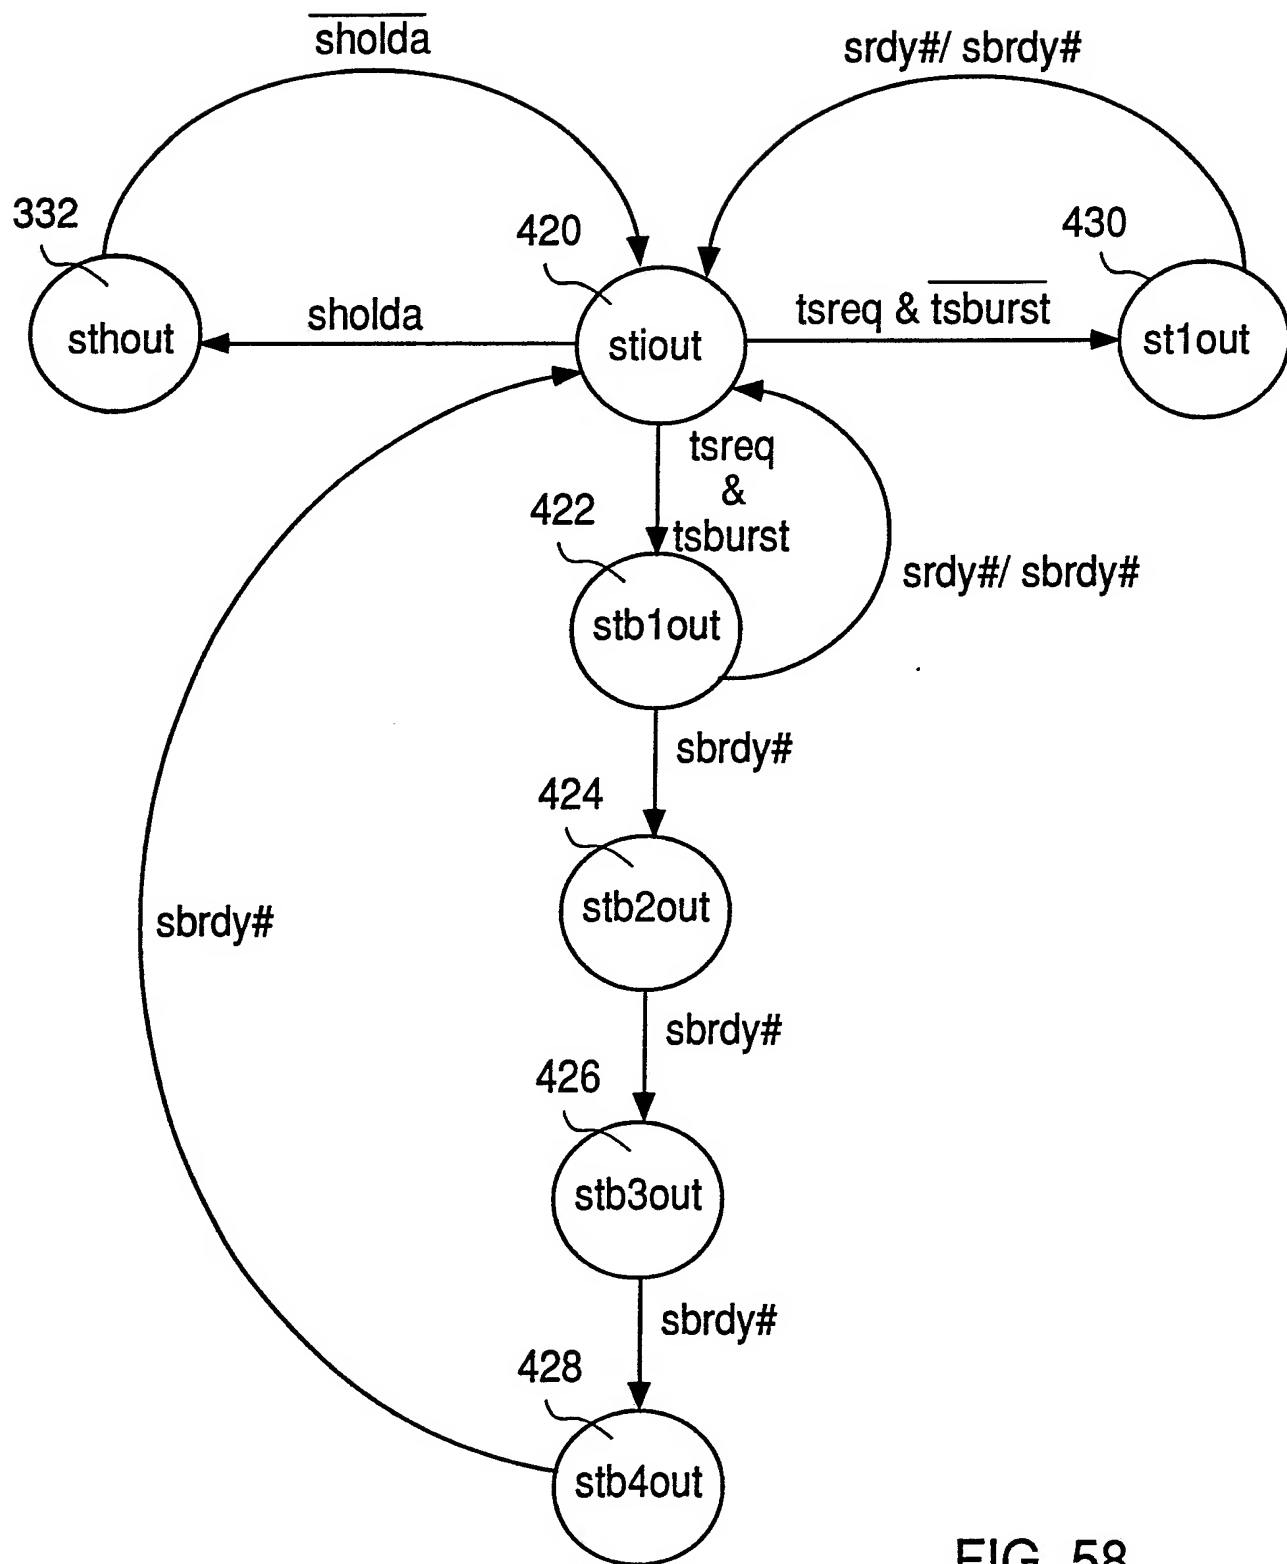


FIG. 58

INTERNATIONAL SEARCH REPORT

International Application No. **PCT/US91/04484**

I. CLASSIFICATION OF SUBJECT MATTER (if several classification symbols apply, indicate all) ⁶ According to International Patent Classification (IPC) or to both National Classification and IPC IPC(5): G11C 7/00 ^{USC} 365/230.05 364/200											
II. FIELDS SEARCHED <div style="text-align: center; border-top: 1px solid black; border-bottom: 1px solid black;">Minimum Documentation Searched ⁷</div> <table style="width: 100%; border-collapse: collapse;"> <tr> <th style="width: 25%; border-bottom: 1px solid black;">Classification System</th> <th style="border-bottom: 1px solid black;">Classification Symbols</th> </tr> <tr> <td style="text-align: center; vertical-align: middle; border-right: 1px solid black; padding: 10px;">US</td> <td style="padding: 10px;">365/49,230.05, 230.08 364/200</td> </tr> </table> <div style="text-align: center; border-top: 1px solid black; border-bottom: 1px solid black;">Documentation Searched other than Minimum Documentation to the Extent that such Documents are Included in the Fields Searched ⁸</div>			Classification System	Classification Symbols	US	365/49,230.05, 230.08 364/200					
Classification System	Classification Symbols										
US	365/49,230.05, 230.08 364/200										
III. DOCUMENTS CONSIDERED TO BE RELEVANT ⁹ <table style="width: 100%; border-collapse: collapse;"> <tr> <th style="width: 10%; border-bottom: 1px solid black;">Category [*]</th> <th style="border-bottom: 1px solid black;">Citation of Document, ¹¹ with indication, where appropriate, of the relevant passages ¹²</th> <th style="border-bottom: 1px solid black;">Relevant to Claim No. ¹³</th> </tr> <tr> <td style="text-align: center; vertical-align: top; border-right: 1px solid black; padding: 10px;">X</td> <td style="padding: 10px;">US, A, 4,403,288 (CHRISTIAN ET AL) 06 SEPTEMBER 1983 See column 8, line 20 to column 9, line 11.</td> <td style="text-align: center; vertical-align: top; border-right: 1px solid black; padding: 10px;">22-29</td> </tr> <tr> <td style="text-align: center; vertical-align: top; border-right: 1px solid black; padding: 10px;">X</td> <td style="padding: 10px;">US, A, 4,577,293 (MATICK ET AL) 18 MARCH 1986 See entire document.</td> <td style="text-align: center; vertical-align: top; border-right: 1px solid black; padding: 10px;">1,2,4,19</td> </tr> </table>			Category [*]	Citation of Document, ¹¹ with indication, where appropriate, of the relevant passages ¹²	Relevant to Claim No. ¹³	X	US, A, 4,403,288 (CHRISTIAN ET AL) 06 SEPTEMBER 1983 See column 8, line 20 to column 9, line 11.	22-29	X	US, A, 4,577,293 (MATICK ET AL) 18 MARCH 1986 See entire document.	1,2,4,19
Category [*]	Citation of Document, ¹¹ with indication, where appropriate, of the relevant passages ¹²	Relevant to Claim No. ¹³									
X	US, A, 4,403,288 (CHRISTIAN ET AL) 06 SEPTEMBER 1983 See column 8, line 20 to column 9, line 11.	22-29									
X	US, A, 4,577,293 (MATICK ET AL) 18 MARCH 1986 See entire document.	1,2,4,19									
<div style="display: flex; justify-content: space-between;"> <div style="width: 48%;"> <p>[*] Special categories of cited documents: ¹⁰</p> <p>"A" document defining the general state of the art which is not considered to be of particular relevance</p> <p>"E" earlier document but published on or after the international filing date</p> <p>"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)</p> <p>"O" document referring to an oral disclosure, use, exhibition or other means</p> <p>"P" document published prior to the international filing date but later than the priority date claimed</p> </div> <div style="width: 48%;"> <p>"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention</p> <p>"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step</p> <p>"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.</p> <p>"&" document member of the same patent family</p> </div> </div>											
IV. CERTIFICATION <table style="width: 100%; border-collapse: collapse;"> <tr> <td style="width: 50%; border-right: 1px solid black; padding: 5px;"> Date of the Actual Completion of the International Search <div style="text-align: center; font-weight: bold;">10 OCTOBER 1991</div> </td> <td style="width: 50%; padding: 5px;"> Date of Mailing of this International Search Report <div style="text-align: center; font-weight: bold; font-size: 1.2em;">04 NOV 1991</div> </td> </tr> <tr> <td style="border-right: 1px solid black; padding: 5px;"> International Searching Authority <div style="text-align: center;">ISA/US</div> </td> <td style="padding: 5px;"> Signature of Authorized Officer <div style="text-align: center;"> NGUYEN NGOC-HO <i>Nguyen</i> INTERNATIONAL DIVISION JOSEPH A. POPEK </div> </td> </tr> </table>			Date of the Actual Completion of the International Search <div style="text-align: center; font-weight: bold;">10 OCTOBER 1991</div>	Date of Mailing of this International Search Report <div style="text-align: center; font-weight: bold; font-size: 1.2em;">04 NOV 1991</div>	International Searching Authority <div style="text-align: center;">ISA/US</div>	Signature of Authorized Officer <div style="text-align: center;"> NGUYEN NGOC-HO <i>Nguyen</i> INTERNATIONAL DIVISION JOSEPH A. POPEK </div>					
Date of the Actual Completion of the International Search <div style="text-align: center; font-weight: bold;">10 OCTOBER 1991</div>	Date of Mailing of this International Search Report <div style="text-align: center; font-weight: bold; font-size: 1.2em;">04 NOV 1991</div>										
International Searching Authority <div style="text-align: center;">ISA/US</div>	Signature of Authorized Officer <div style="text-align: center;"> NGUYEN NGOC-HO <i>Nguyen</i> INTERNATIONAL DIVISION JOSEPH A. POPEK </div>										